Wide-Bandwidth High-Resolution Search for Extraterrestrial Intelligence

Semiannual Status Report

15 June 1993 -- 15 Dec 1993

Paul Horowitz, Principal Investigator

Harvard University
Cambridge, MA
02138
December 15, 1993

Grant Number NAGW-2872

1. INTRODUCTION

This report summarizes research accomplished during the fourth 6-month period of the grant. During the period covered by this report the active personnel included the PI, two Harvard graduate students (Darren Leigh and Jonathan Weintroub) and an MIT graduate student (Max Avruch), each averaging a 75% research commitment; a Harvard undergraduate (Derrick Bass), and a recent mathematics graduate from Harvard (Nick Shectman).

2. RESEARCH ACCOMPLISHED

2.1 Antenna System

As described previously, the dual feedhorn required for our two-lobe sky pattern is complete and tested, as are the amplifiers, downconverter, mixers, and digitizers -- all the analog "upstream" components. Early in the project, however, we decided to add a third antenna -- a terrestrial low-gain feed, to act as a veto for local interference.

This terrestrial antenna has now been designed, built and tested. We needed a wideband (about an octave) antenna with an azimuthally symmetric pattern. We decided to use a discone design. The discone antenna is vertically polarized, has a symmetrical power pattern and an extremely wide bandwidth. It also has the added advantage of filtering out almost all energy at frequencies below a certain cutoff. We decided to use a lower cutoff frequency of 1 GHz and a usable frequency range of 1.2 GHz and up. At just 6 cm tall, this is probably the smallest discone antenna ever built. During testing we verified that it had a transmitting return loss of better than 10dB over the entire frequency range of interest. This is more than adequate for our needs.

2.2 FFT Array

In previous reports we have described our innovative 3-chip design for a 4 megapoint complex FFT, wherein the long transform is implemented as a succession of shorter row and column FFTs, with complex ("twiddle") multiplies interposed between the shorter transforms during the "corner turns" (Figure 1). This architecture builds on the elegant Serendip design of the Berkeley SETI group. We also described numerical simulations, verifying detection of weak narrowband signals in the presence of much stronger signals and broadband noise (see the 12/92 progress report for an exhaustive description). Those simulations established the degree of window- and twiddle-ROM truncation that can be used (to reduce cost), consistent with desired dynamic range and spur-free response.

2.2.1 Hardware Implementation and Features

Our major accomplishment during the period of this report is the successful reduction of the FFT architecture to finished working hardware. Figure 2 shows the final block diagram. The data, initially quantized to 8 bits in both I and Q, passes through a 4M bit-reversing "corner turn," then through a 16K FFT (at 20-bit precision) implemented as 128x128 points (with an input windowing ROM, an internal 16K corner turn, and an internal 16K twiddle multiplication). After another 4M corner turn the data passes through a 256-point FFT, this time with a 4M

twiddle multiplication. The output spectrum (20-bit complex I,Q amplitudes, in bit-reversed order) finally is converted to 16-bit moduli by a 2-step ROM-lookup with input "p-code" compression (described fully in the 6/93 progress report), then bit-reversed in a final 4M DRAM corner turn.

There are several unusual aspects to our implementation. First, we found a way to eliminate a second 16K corner turn that follows the 16K composite FFT, by absorbing its 14-bit permute into the following 22-bit permute. This simplifies both hardware and timing, but is not entirely simple: The concatenated permutes do not display the cycle-of-two periodicity that a simple end-to-end bit reverse enjoys, thus cannot be simply implemented in unbanked (read-modify-write) memory, as are the other three corner turns. Rather, the combined permutation displays a cycle-of-12! As with the other address manipulations, the resulting complex addressing pattern can be implemented in a complex PLD ("CPLD," in this case an AMD Mach-220). Furthermore, the resulting 22-bit logical address can be mapped onto DRAM physical rows and columns in such a way that the required DRAM refresh (access of all combinations of the bottom 10 row bits in no more than 16 milliseconds) is automatically accomplished during the FFT; by actual calculation, the maximum row address revisit time is 9.9 milliseconds over the full cycle-of-12.

Another useful feature is the collection of error flags into a status word that is substituted for the first word of the 4M output vector: During the computation of the FFT, parity is checked at the input port, and at each of the 9 DRAM SIMMS; these ten error bits are held in 1's-catching registers, along with the three overflow error bits of the FFT chips. The resulting 13 bits comprise the error status word. If any of the 13 bits are non-zero, the corresponding spectrum will be ignored; persistent errors, of course, signal faulty hardware, and require corrective maintenance.

To aid in maintenance, there are a set of headers for monitoring the passage of data through the FFT (shown at the top of Figure 2), along with corresponding timing signals. To use this feature, we substitute a digital pattern generator (a one-chip circuit using a Mach-110) for the data input, then capture data with a logic analyzer at each successive header. We have simulated (and checked against working hardware) the data streams at each header. One simply compares these normal "signatures" with the actual data flow, thus isolating the hardware error between a pair of headers. By examining the errors, we may be able to identify the faulty hardware precisely (e.g., a stuck bit in a register). Even if we are not that lucky, at least we will have isolated the problem to at most 3 or 4 chips, which can then be substituted (all chips are socketed) until the problem is fixed.

As yet another aid to reliable operation of this very large spectrometer, we have included a pair of addressable test ports at input and output. In "test" mode (implemented by setting bit 1 of chip 3 through the supervisory port) these ports substitute for the normal data ports, permitting an external data source to drive test vectors through the FFT array. In practice we plan to send identical data (even antenna noise will do) through a chosen pair of FFT cards, comparing the results for identity; by cycling through the cards we can verify correct operation of the full array of 63 cards in approximately one minute.

Note that the data flows only through registers and signal processing chips, while control and addressing is generated by a set of four synchronized CPLDs.

2.2.2 4-Megachannel Printed Circuit Board

Figure 3 shows the complete schematic of the 4M FFT board, in an (unreadable!) "flat" single-page circuit. It was then hand placed and routed to create the artwork shown in Figure 4, the two signal layers of a 4-layer PCB. We used 8/8 design rules, with "double-track" routing.

Figure 5 shows the completed board. It contains 36 MByte of DRAM (the 9 SIMMS along the bottom), 5 CPLDs (in PLCC sockets), the three large FFT ASICs (in ceramic 132-pin PGAs), and lots of "glue." There are 74 ICs in all, with a total parts cost of about $2000.

It is interesting to compare this board with the META array of 144 circuit boards, each about 25% larger than this 4M BETA board. The BETA board has half the channel count of the full META array, but ten times the channel bandwidth. A reasonable figure of merit is channel count times channel bandwidth. Thus a single BETA board is five times as powerful as the full META array!

2.2.3 Performance

It took close to a week to fully debug the 4M board. The only problems were a half-dozen errors in the CPLD boolean equations. We fixed these "firmware" errors, and did not have to make any changes at all in the printed circuit hardware. To speed debugging, we built a board (Figure 6) that grabs 256-point or 4K-point "snapshots," beginning at a designated point in the spectrum, and forwards them to a vector analog display card with data buffer (Figure 7); the latter was built to debug the Arecibo Spectrometer/Power-Accumulator described below. An embedded microcontroller on the snapshot board controls the process, as well as permitting flexible live-time configuration of the 4M FFT board's control registers. We have also constructed an interface to a PC, so that these 256-point snapshots can be logged to disk for subsequent analysis and plotting.

Figures 8 through 11 show the 4M FFT board's performance when driven by a variety of input signals. We plotted these graphs by feeding the PC interface with the snapshot buffer. Each graph is a plot of modulus (from the board's 16-bit output words) versus channel number, for 256 points grabbed from the appropriate starting point within the 4M output points.

Figure 8 demonstrates the response to a weak sinusoid embedded in wideband noise. For a noise source we used cascaded IF amplifiers (Avantek GPD-type) to amplify resistor Johnson noise, with a bandpass filter (at 70 MHz) and further amplification to preserve headroom. We combined the noise with a VHF synthesizer, locked to the 40 MHz clock that times the FFT. Figure 8a shows the response to a sinewave centered on a frequency bin (the carrier was set precisely 0.5 MHz above the local oscillator [LO]), with no windowing. This is an artificial situation, as most signals will not have that lucky alignment; but it does illustrate the ability to detect, at high signal/noise ratio, a signal that is 40dB below noise in the analysis bandwidth.

Figure 8b shows a more typical situation, with the carrier displaced 0.4 bins from bin center. Note the reduced peak response ("picket fence" effect), as well as the spreading of the signal to several adjacent bins (spectral "leakage"). Leakage is particularly serious when strong signals are present (see below), and is generally handled by "windowing" -- multiplying the input time series by a smooth positive function that goes to zero at the ends of the analyzed time series. (In optics the spatial analog of this is known as "apodizing".) Figure 8c shows the result with a Hanning window ("von Hann," henceforth VH; note expanded vertical scale), and Figure 8d shows the effect of a Blackman-Harris window (henceforth BH). Windows are not terribly important with a weak signal in wideband noise, but their utility will become apparent below.

A 4-million-point FFT should enhance the S/N ratio of an unresolved carrier in broadband noise by a factor of 4M, or 66dB. Thus the carrier here, at a representative offset from bin center, should be seen at +26dB (a factor of 400 in power, or 20 in modulus) above rms noise in Figures 8c,d (and somewhat better in 8a), as it is.

Figure 9 shows the more interesting 2-tone test: Here we've put a strong carrier exactly midway between bins ("mid-bin"), and a much weaker carrier (40dB down) just 10 Hz away. At this resolution that's 21 bins away, again mid-bin. In Figures 9a and 9b we've done the FFT unwindowed (euphemistically called a "rectangular" window), displaying the serious spectral "leakage" resulting from the convolution of a *sinc* function in the spectral domain (the FFT of a rectangular impulse). Because the weak sine combines coherently with the spectral leakage tail of the stronger signal, the rotation of their relative phases produces a non-stationary combined amplitude, illustrated in Figures 9a (best relative phase) and 9b (worst relative phase -- nearly complete cancellation of the weaker signal). In the latter, the weak sine is lost in the leakage wings.

In Figure 9c we've used the Hanning (VH) window, and in Figure 9d the BH window. The BH is a severe window, with peak sidelobes of -92dBc, bought at the expense of a fairly broad main lobe (a pure sinusoid typically becomes 4 to 5 channels wide in the frequency domain). The VH has a peak sidelobe level of -32dBc, but then falls at 18dB/octave; its main lobe is typically 3 channels wide. These characteristics are apparent in Figures 9c,d: Note the narrower top portion of the main peak, when using the VH window, but broader width of the base (this is a logarithmic plot), compared with the BH window. Both satisfactorily separate the weaker signal from the sidelobes of the stronger; the comparison with the rectangular window is stunning.

Figure 10 shows similar spectral experiments, this time with a weakly modulated AM signal. Here the carrier, at an IF frequency of 70.5 MHz, is square-wave modulated at 5 Hz, and mixed with a 70 MHz LO. Once again, well-constructed windows are the salvation of spectroscopy.

Some people think that a short tone burst can produce a narrow spectral line, thus mimicking the unresolved carrier we seek. In Figure 11 we've obligingly analyzed tone bursts: Beginning with Figure 11a (continuous carrier, present for the full 2-second time window of the FFT), we successively shorten the tone duration by factors of 2. Thus Figure 11b is a 1-second burst, 11c is a half-second burst, etc. The resulting spectral broadening, with *sinc* envelope, is readily apparent; so is the reduction in peak amplitude.

2.2.4 Production

We are most pleased with the performance of the 4M FFT spectrometer board, and will move forward with a production run of approximately 70 cards in the next month. Because of a budget shortfall in 1994 (due to congressional termination of NASA SETI), we will not be able to stuff the full set of boards in the next grant period. Instead, we will produce and field finished boards in a fund-limited mode.

2.3 Radioastronomy Spectrometer

In the 6/93 progress report we described the design of a combined spectrometer and power-accumulator ("SPA"), for use at Arecibo Observatory to search for neutral hydrogen emission from condensations at high redshift (z=5). This project, a technological spinoff from our SETI work, is progressing well. Our proposal to use the Arecibo radiotelescope for the search has been approved by NAIC. The work is supported by non-NASA funding, and is a collaboration with Bernie Burke of MIT, Mike Davis of NAIC, and Jim Cordes of Cornell, who will be using the spectrometer for pulsar studies.

The Austek FDP-based 256 point Spectrometer/Power Accumulator (SPA) has now been implemented as a 4-layer printed circuit, using relaxed 12/12 design rules, and five prototype boards have been fabricated, of which four have been assembled; see Figure 12. The handwired vector display tester board mentioned earlier (designed, built and debugged during the period covered by this report) interfaces to the digital output of the SPA and presents the integrated spectrum as an X-Y display on an oscilloscope. The SPA board has been tested with deterministic input vectors and the output compares point by point with that obtained using the Austek simulator running on our Sun workstation -- a convincing demonstration of correct operation.

More interesting, though, is the board's ability to operate on "real" signals. We tested performance in the laboratory using both narrow- and broad-band signals buried in noise and demonstrated convincingly that the process of integrating spectra properly enhances the signal/noise ratio of signals buried in noise. As with the SETI spectrometer, the SPA spectrometer operates on inputs that are filtered, mixed to IF, and digitized using the "Mixer/Digitizer" boards originally designed for SETI. Sample spectra are shown in Figures 13 and 14. In Figure 13 a weak carrier (-30dBc in the analysis bandwidth) is shown after 512, 8K, and 128K integrations; the S/N ratio grows as the square root of the (incoherent) integration time, thus quadrupling from each plot the next. The dip at DC is due to AC coupling at the ADC input, with a -3dB point of about 2kHz, which suppresses the DC channel and one channel to either side.

In Figure 14 we have displayed similar processing for a flat-topped spectral line (produced by triangle-wave FM), whose 100kHz bandwidth is typical of what we might see from doppler-broadened neutral hydrogen. Once again, the improvement obtained by incoherent integration is impressive. The spectra presented in these hard copy plots were captured using the PC computer interface discussed next.

A PC interface board has been designed and two units built and tested. The interface board and its associated user interface and control software allows an IBM compatible computer to control the SPA board, and facilitates the transfer of spectra to the PC for display, processing and storage. One of the interface boards has already been shipped with its software, SPA board and Mixer/Digitizer boards to the Arecibo observatory. Mike Davis has used this equipment, with Arecibo's L-band line-feed, to observe actual 21-cm galactic hydrogen. The spectral lines walked through the display, as the beam transited the galactic arm structure.

Members of our group and our MIT collaborators will travel to Arecibo in January 1994, at which point we hope to have two (east-west) low frequency helical feeds mounted to the Arecibo catwalk. This will allow us to bring a scaled-down version of the early hydrogen spectrometer on-line in the 200-300MHz band. The production run of spectrometer printed circuit boards will be ordered shortly. As SPA boards are assembled and tested in Cambridge we will ship them to Arecibo and add them to the system; we expect to be at full bandwidth by June 1994. We have also decided to build our own IF filter bank and LO array, thereby spinning off the SETI designs described in earlier reports. We consider that being able to operate completely independently of shared Arecibo electronics (such as the filter bank) is well worth the additional trouble and expense.

One of the five prototype SPA boards has been sent to Jim Cordes at Cornell. He has started work on the much higher speed computer interface needed for pulsar work, and has also found application for the spectrometer and PC interface in the teaching program at Cornell (a SETI spinoff into education). We continue to work closely with him.

2.4 Feature Recognizer

We have added significant enhancements to the design of the Feature Recognizer Array. These cards receive the stream of modulus words from the 4M FFT cards, and forward a greatly thinned set of reports to the PCs in whose backplane they reside. In particular, we have adopted a powerful ROM-based state-machine architecture, and we have added DRAM to permit integration modes when tracking or reobserving source candidates.

In the current architecture there are 21 '486-class machines, each with three Feature Recognizer boards, one Feature Correlator board, and one ethernet board. The Feature Recognizer boards each talk to one Spectrometer board and to the Feature Correlator board in the PC, and the Feature Correlator board takes data from the Feature Recognizers, and decides, based on information about that data that it receives from the Feature Recognizers and from the PC, whether to pass the data on to the PC for further reduction. This design reduces the number of PC interfaces required per set of boards from three (in the previous design) to one, and simplifies other circuitry. It also provides greater flexibility for changes in operation after the boards have been manufactured, by allowing for functional changes by reprogramming the PALs and EEPROMs described below. In addition, we have added DRAM to the system in a way that allows the system to operate with the DRAM only partly populated (or not at all), allowing the system to go on line at a low cost and be upgraded later, when DRAM prices are lower, if we decide we want the added functionality.

The design of the Feature Recognizer board is nearly complete, and the Feature Correlator is at the block diagram level. In addition, a port-level specification of the PC interface has been completed, allowing work on the back-end software to progress. We expect to complete both designs and debug the prototypes during the next grant period.

2.4.1 Feature Recognizer Operation

The Feature Recognizer boards have four modes of operation: regular operation (trolling), integration, and two readout modes, one in which the integrated magnitude in selected regions is read out through the post-processing features used during regular operation, and one in which the full power spectrum is read out directly, to be written to disk or tape and reduced later.

In regular operation, the Feature Recognizer board keeps a running average of 4K data points (root mean of squares) and compares each data point simultaneously to four thresholds computed by multiplying the root mean of squares of local (2K on either side) magnitude by four programmable multipliers. The Feature Recognizer presents the data and the average data to a 32-bit bus, and also provides the results of the comparisons for the Feature Correlator to use in correlating data and deciding whether to send it to the PC. In addition, the Feature Recognizer attached to the east horn also provides some timing signals and an address for the data to the Feature Correlator, and makes its DRAM available to the Feature Correlator for use by the Feature Correlator's state machine.

In integration operation, the Feature Recognizer integrates (in power) 4M spectra in its DRAM. A special mode is used for the first spectrum of any integration to zero out the DRAM.

In reprocessed readout mode, the integrated power is converted back to RMS average amplitude (averaged over time, not over frequency) and fed through the feature recognizing circuitry exactly as if it were incoming data coming from an FFT board, and presented to the Feature Correlator in the same way. This mode allows the Feature Correlator to select those portions of the integrated spectrum that it thinks are interesting, to make up for the fact that the PC's data bus is too slow to allow a full readout of a 4-Mpoint integrated spectrum in one 2-second cycle.

In direct readout mode, the integrated power is not converted to amplitude, but is read directly onto the bus to the Feature Correlator, with no advisory information. This process must be repeated several times, with a different section of the spectrum sent to the PC each time, until the entire spectrum has been read out; that is because the PC's bus cannot sustain the high data rate from the FFT.

2.4.2 Feature Correlator Operation

The core of the Feature Correlator consists of a state machine, implemented as a fast EEPROM or SRAM. There are also two FIFOs, one for storing "slot" and "notch" information from the PC, to be used to force or suppress data reporting, and one for storing actual data being reported to the PC. Finally, there are two groups of PALs, one for controlling the output FIFO and driving the synchronous parts of the Feature Correlator and the synchronous signals on the

connectors to the Feature Recognizer, and one for controlling the PC interface and all asynchronous signals, and for reprogramming the state machine.

In regular operation, the state machine considers encoded feature information from each of the three Feature Recognizers, advisory information from the PC, and state from the previous frequency bin, held in a latch, and from the same frequency bin from the previous spectrum (actually, since the main local oscillator "hops" through 8 values to increase the overall bandwidth of the system, it uses the state from the spectrum 8 cycles previous), held in the DRAM of one of the Feature Recognizer boards. It generates new state for the latch and the DRAM, and produces an opinion as to whether the data should be reported. This opinion is combined with "slot" and "notch" directives from the PC, and if the final decision is to report the data, the data and baseline information from all three boards, as well as the address of the data and a set of flags describing the information, are written to a 128-bit-wide FIFO, which can be accessed as 8x16-bit I/O ports on the PC.

In readout operation, the Feature Correlator sifts through the readout data being provided to it by the Feature Recognizer boards, and either reports "interesting" data to the PC, using the same criteria that it uses in regular operation (except for DRAM-based state), or alternately uses a "comb" pattern of slots and notches to read out an entire integrated spectrum over several 2-second cycles. The Feature Correlator can also use slots and notches, or the comb pattern, to pass direct readout data to the PC.

The most interesting feature of the Feature Correlator is its ability to operate the different Feature Recognizers in different modes. One promising possibility is to operate the east horn in regular mode, with "hopping", to scan for interesting features with rejection from the terrestrial horn, but to operate the west horn with a non-hopping, but still programmable, LO, to do immediate followup of interesting signals in integration mode, using slots and notches to look at the interesting parts of the spectrum in direct readout mode. Another possibility is to populate the DRAM in two horns and run both in integration mode, adding the results of an east integration to the results of a slightly later west integration as the telescope trolls. The Feature Correlator's PALs can be reprogrammed to support new modes as we devise them.

In addition to having 16 ports mapped into the PC's I/O space, the Feature Correlator will also map its state machine into memory space in program operation, so that changes in programming of the state machine (but not the PALs) can be done remotely, without removing the cards from the system. In fact, since the Feature Correlator's PC interface is unused during Feature Recognizer integration mode, state machine reprogramming could happen over the course of an integration, without losing telescope time.

2.4.3 Feature Recognizer Summary

The current Feature Extractor design improves the original rigid hardware design (which was cheap and fast, but could only do fixed-protocol hit/slot/notch processing) by using a microcoded state-machine architecture to support many data-sifting modes. The most interesting of these use memory to store a state bit or two (to force reduction of subsequent thresholds), or to perform true power-law digital integration to enhance detected S/N ratio.

Thus our design wins back some of the flexibility of general-purpose processors, and adds integrating modes and state memory, while maintaining high speed and low cost, in a memory-expandable implementation.

2.5 General Purpose Array and Workstation

The general purpose (GP) array consists of twenty '486 PC class computers, each of which receives and processes the data from a feature extractor/correlator board set. The array performs a first analysis on the provided "features" and then passes this information on to the workstation. Each computer in the array is responsible for real-time communication with both the feature extractor board set and the ethernet. In order for the array's software to be well organized as well as fast, it is being written in C++ using a DOS extender which provides access to the Intel '486 microprocessor in its native mode. This not only provides us with the full speed of the '486, but also allows us to use a "flat" memory model (no segments or 640K barrier) which makes the job of dealing with large amounts of data much easier. Since high speed and low latency are important requirements, the array software deals intimately with the '486 motherboard interrupts. The C++ programs are linked with hand crafted assembly language subroutines which handle these. The result is a fast, clean interface between the high level programming code and the motherboard hardware. This allows the programmer to use high level, object-oriented constructs to access the needed hardware features without having to worry about timing or resource management issues; they have already been taken care of. This code has been exhaustively tested via communication between a PC and the workstation.

The array computers receive their programs over the network in a diskless fashion. In order to do this the ethernet adapter card in each computer requires a special boot ROM. The assembly code programs that go in the ROM have been finished, debugged and extensively tested. The network boot function works well with all of our software, including programs running under the DOS extender which perform direct hardware I/O.

We are currently working on the "feature" preprocessing software and finishing up the PC-to-workstation communication routines. The array PCs will be operating unattended for long time periods using finite resources to run programs with many data dependencies. Because of this, we are using special resource management techniques to make sure that the PCs do not suddenly run out of memory or get caught in a deadlock condition. This involves keeping requests and pending "tasks" in prioritized queues, using only static storage or allocating memory in fixed size, specially tracked pieces, etc. The queuing and resource management software is finished and interfaces well with object oriented programming constructs.

Progress has been made on pinning down the flexible boundary between the backend workstation and the PC array. The workstation will communicate with the PCs via a standard coax "thin" ethernet. Our experiments have shown that data transmission errors on the large and busy network that serves the Harvard physics department occur about every 500,000 packets. Since our SETI-dedicated network will be better behaved, we have decided to use a raw ethernet protocol. We have defined the data that will be passed between the workstation and the PC, and the protocols that will be used to send them.

The general backend communication architecture is a set of 20 PCs (each connected to its set of spectrometers) which will communicate across the ethernet only when the workstation requests them to. There will also be a "real-time czar PC" that will perform such tasks as maintaining an accurate timebase for the system. It will inform the computers on the network of the time, the local oscillator frequency, antenna position, polarization and so on. It has the right to broadcast at any time and will thus be the only source of packet collisions. These are well handled by the ethernet hardware itself. This czar PC will also control any hardware in the experiment that needs low latency communication.

The core workstation software is now written. That is, the communication channels between the user interface, the backend monitor program and the PCs have working software. We must work next on data analysis and archiving. Recognizing extraterrestrial signals in such a rich data set offers many possibilities for analysis. We are exploring standard data fitting algorithms, neural networks and genetic algorithms for this purpose. A useful user interface still needs to be written.

3. NEXT STEPS

We have made major progress during this 6-month period, most significantly on the complete and debugged 4M FFT board (and its spinoff radioastronomy spectrometer). With all upstream systems now complete (antenna, amplifiers, downconverter, channelizer, and mixer/digitizer), the remaining modules are the Feature Recognizer hardware, and the General Purpose Array and Workstation. The latter are difficult and important; we think they are progressing well. In parallel we will be producing the 4M FFT boards (and radioastronomy spectrometers) in quantity, and knitting the system together in racks.

The budget shortfall will delay fielding of the complete system; however, we expect to be able to put a partial system on-line during the summer months of 1994. Overall, we are very pleased with recent progress.

4. OTHER FUNDING

During the period of this report we have received funding from The Planetary Society, the Bosack/Kruger Charitable Foundation, and Dr. John Kraus, in addition to our grant of partial support from NASA. The Radioastronomy projects are separate from the NASA activities, with support from the two private sources above, and additional funding from NSF and NAIC.

5. PUBLICATIONS AND TALKS

Our paper, "Five Years of Project META: An All-Sky Narrowband Search for Extraterrestrial Intelligence," by Horowitz and Sagan, was published in *The Astrophysical Journal* in September (**415**, 218-235, 1993). A copy is appended to this report. Also published during September was a less technical article on the same subject, "Project META: What Have We Found?" (The Planetary Report, **13**, 5, 4-9 [1993]).

The Principal Investigator gave a talk at Boston's Museum of Science, sponsored by the Wright Center for Science Education at Tufts University. The session was titled "The Onset of Culture

and Prospects for Alien Life," and was organized in an unusual way: Each speaker spent 30 minutes describing his subject; then we moved to comfortable seating around a table with the moderator (Eric Chaisson), at which we responded to audience questions and to each other. By all accounts it was a highly successful evening. He also wrote a short piece that featured student participation in SETI for Odyssey magazine, a publication for school children. In additional school education efforts, Jonathan Weintroub spoke on SETI to 4th graders at the Cambridgeport Elementary school; and Darren Leigh spoke on SETI to minority students at the Charlestown Boys and Girls Club.

We continue to enjoy media coverage, variously in newspapers, magazines, radio, and television. Most recently we were filmed for a half-hour SETI feature on "Future Quest," a PBS series airing in the fall of 1994; the recent *Nova* broadcast featuring the highlights of its 20 years on the air included a segment on our META search project (narrated by Bill Cosby).

6. ACKNOWLEDGEMENTS

We are grateful to Bernie Burke (MIT), Jim Cordes (NAIC Cornell), Michael Davis (Arecibo Observatory), Jon Hagen (NAIC), Dan Werthimer (UCB), and Bill Yerazunis (DEC) for valuable discussions. The 4-megapoint spectrometer grew out of the Serendip design developed by the SETI group at the University of California, Berkeley.

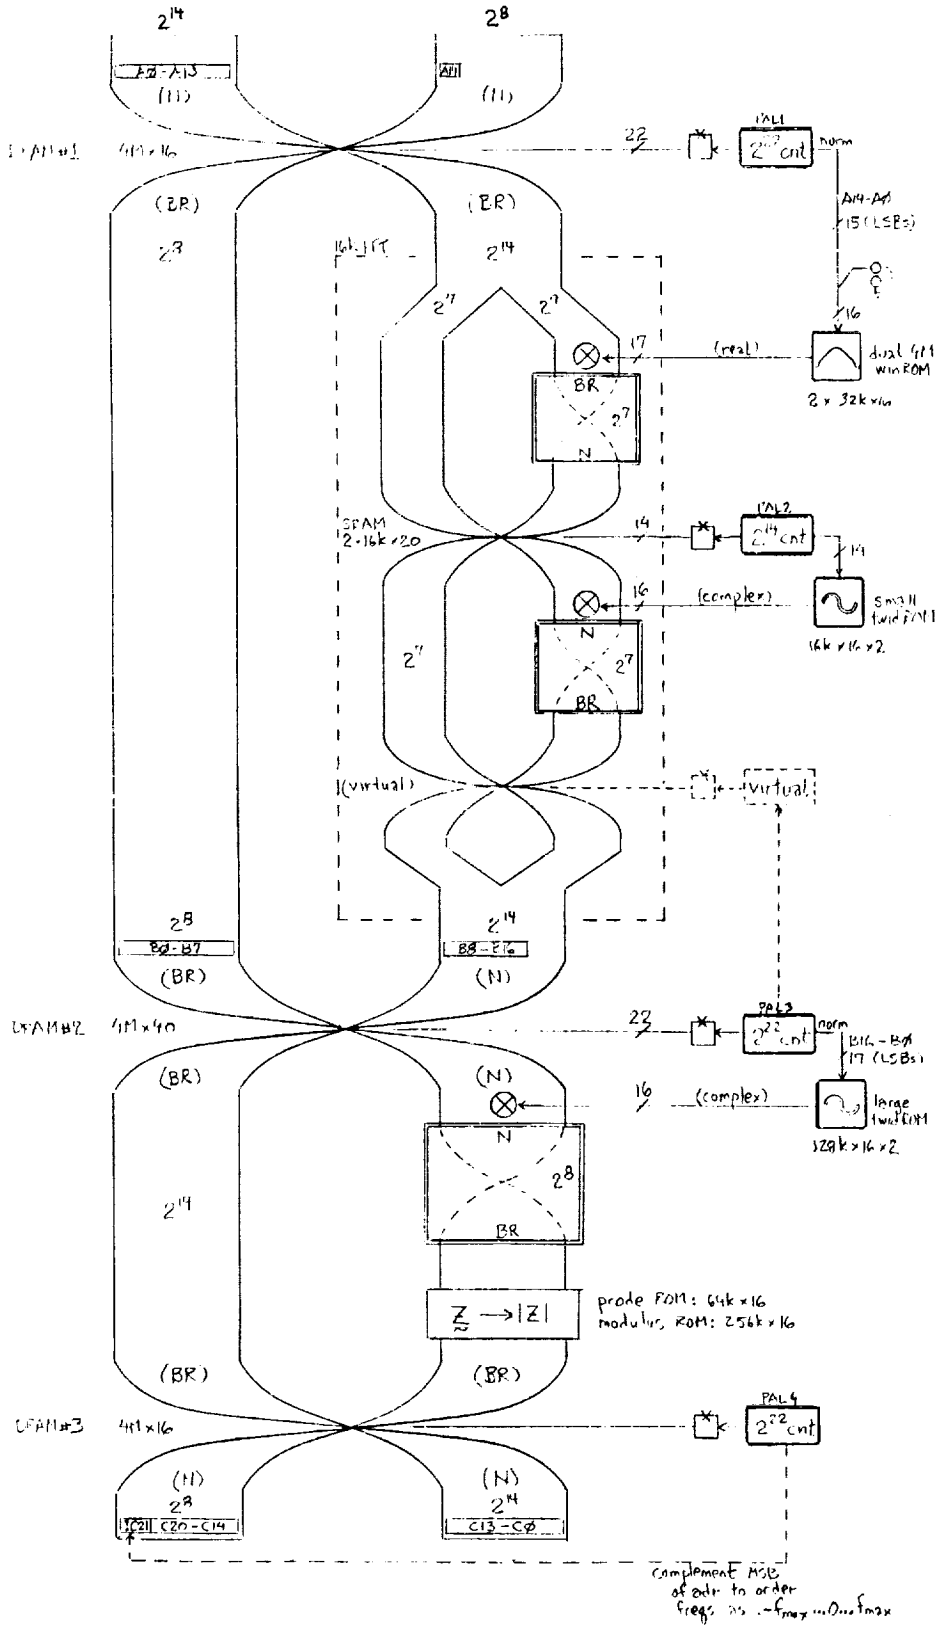

Figure 1. Data flow for a 4-megapoint one-dimensional complex FFT, implemented as a 16K x 256 point 2-dimensional transform with corner turns and complex twiddle factor multiplication. Careful choice of data direction in the small FFTs permits the use of bit-reverse address permutations, rather than conventional transposes. The 16K FFT itself is implemented as 128 x 128 points, again with bit-reverse corner turn and twiddle factors. The three large (4M) corner turns are implemented as read-modify-write cycles in DRAM; the last one is used to produce naturally-ordered output, needed for single-pass baseline averaging in the subsequent signal detector. The multiplier in the first FFT chip is used for time-series windowing.

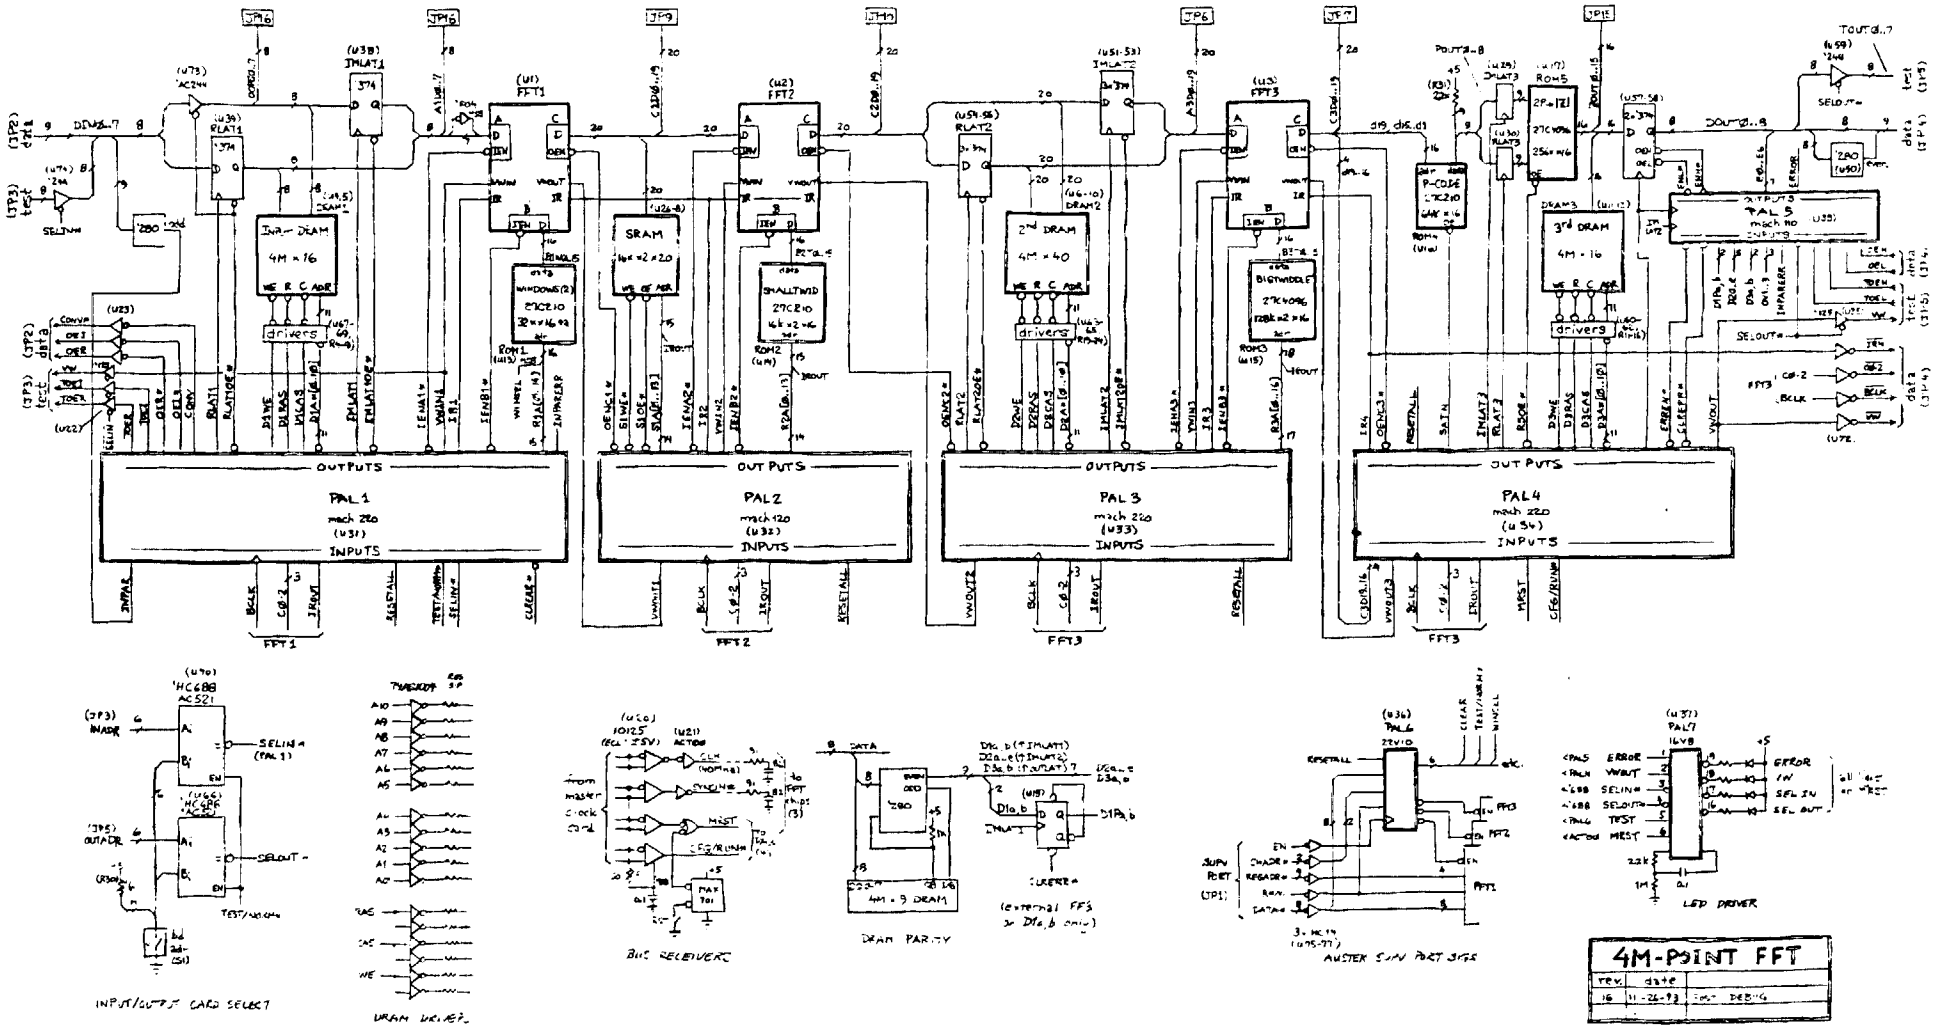

Figure 2. Block diagram of the 4-megapoint FFT. Complex input data, quantized to 8 bits in I and Q, is processed at 20-bit integer precision, and finally converted to a 16-bit unsigned modulus stream. Complex PLDs perform all addressing, timing, and control functions. This circuit computes a pipelined 4M FFT every 2 seconds. Addressable test ports permit comparison of any pair of cards, for system verification. Parity is checked at the input and at each DRAM, and the accumulated error status is substituted for the first spectral point.

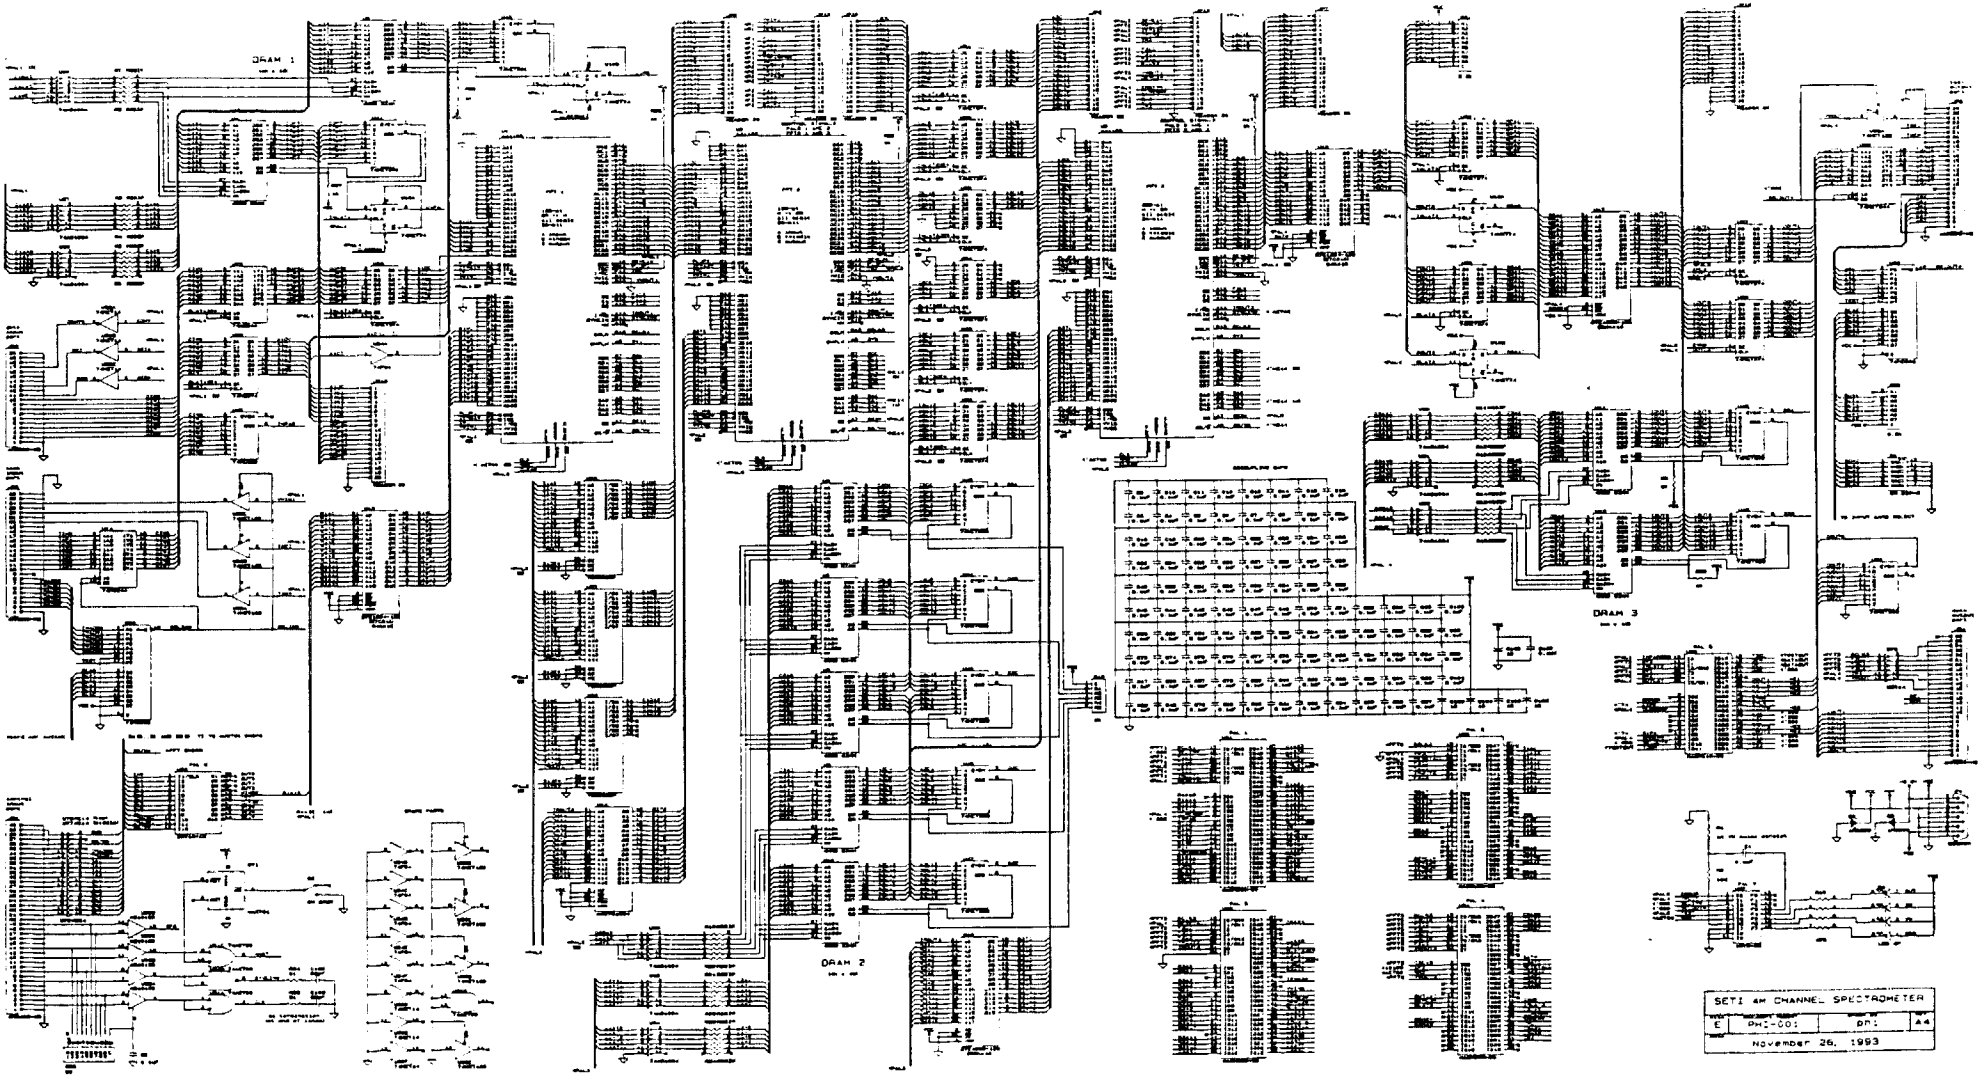

Figure 3. Complete schematic of the 4-megapoint FFT.

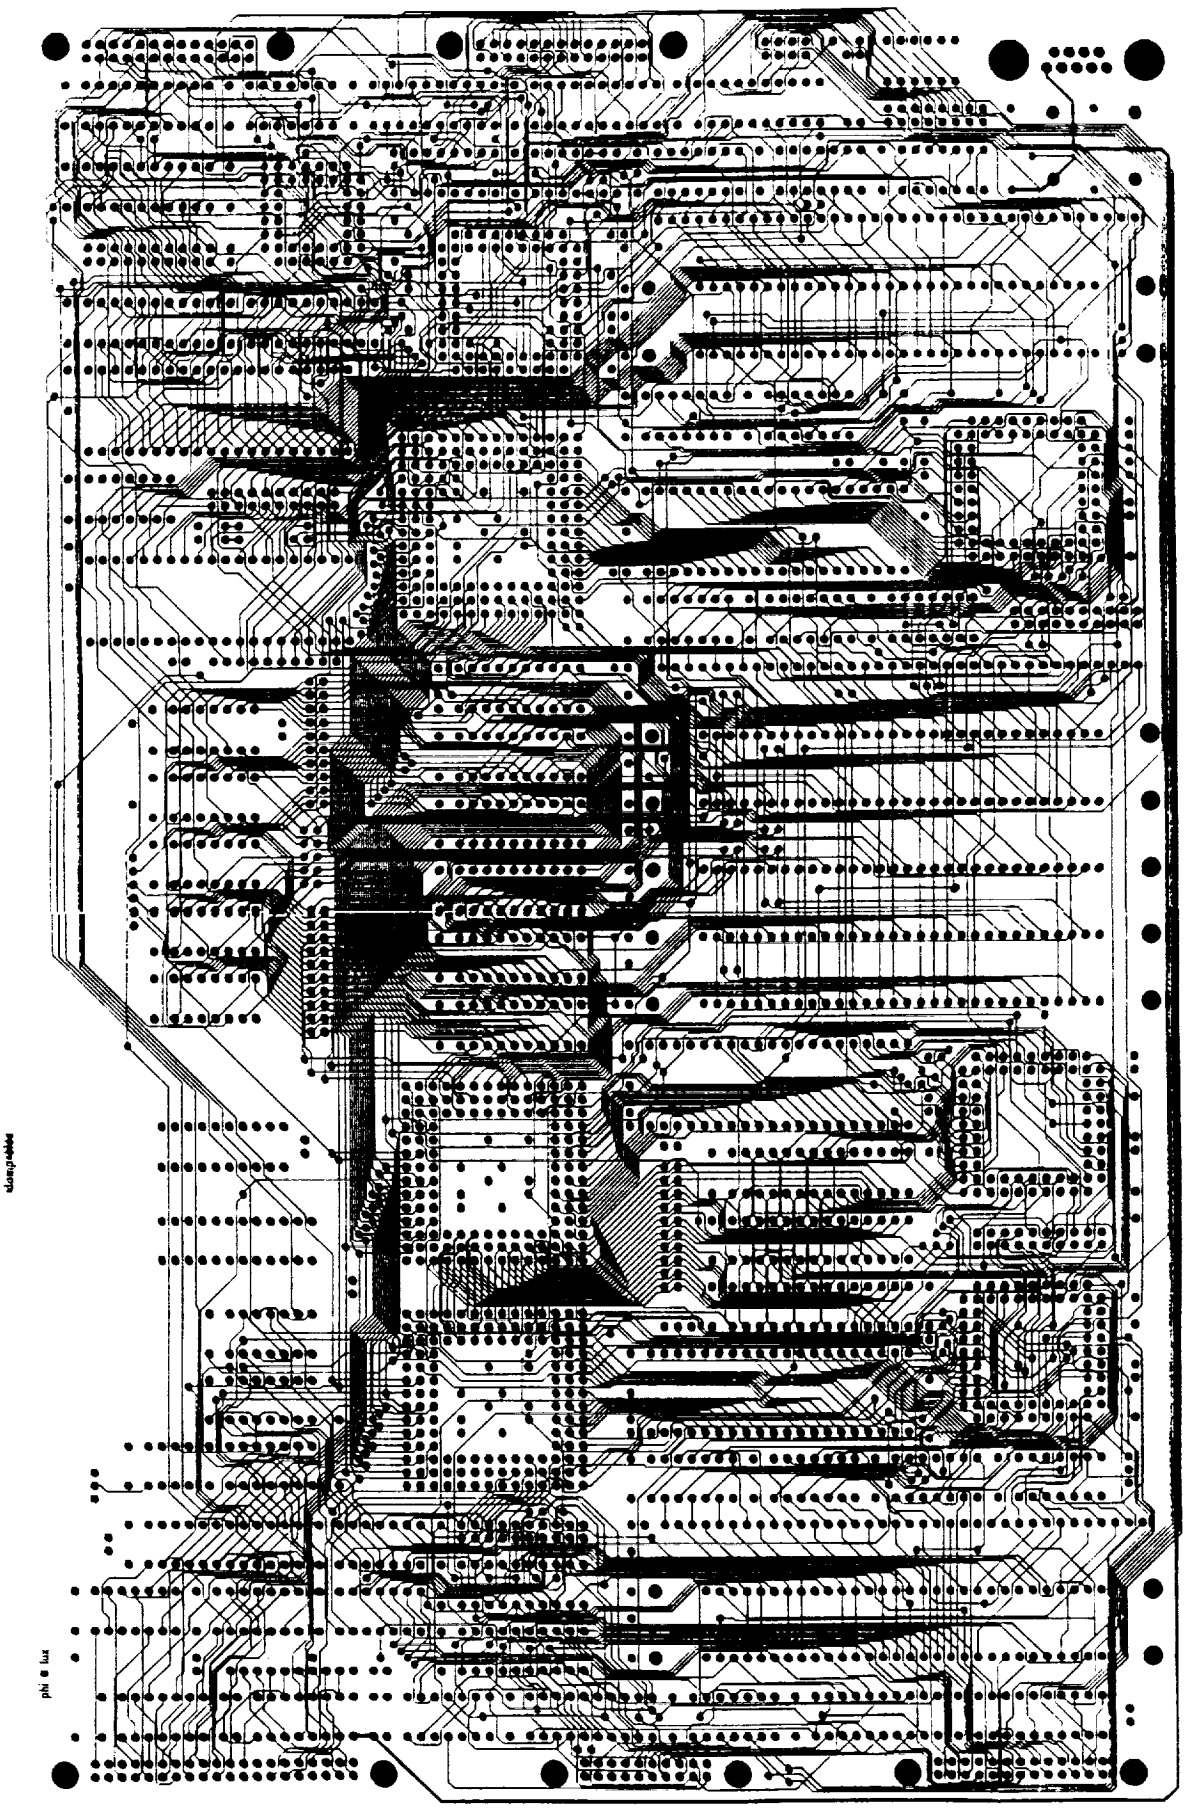

Figure 4. Signal layer artwork for the 4-megapoint FFT printed circuit board. The 4-layer board uses 8/8 double-track design rules.

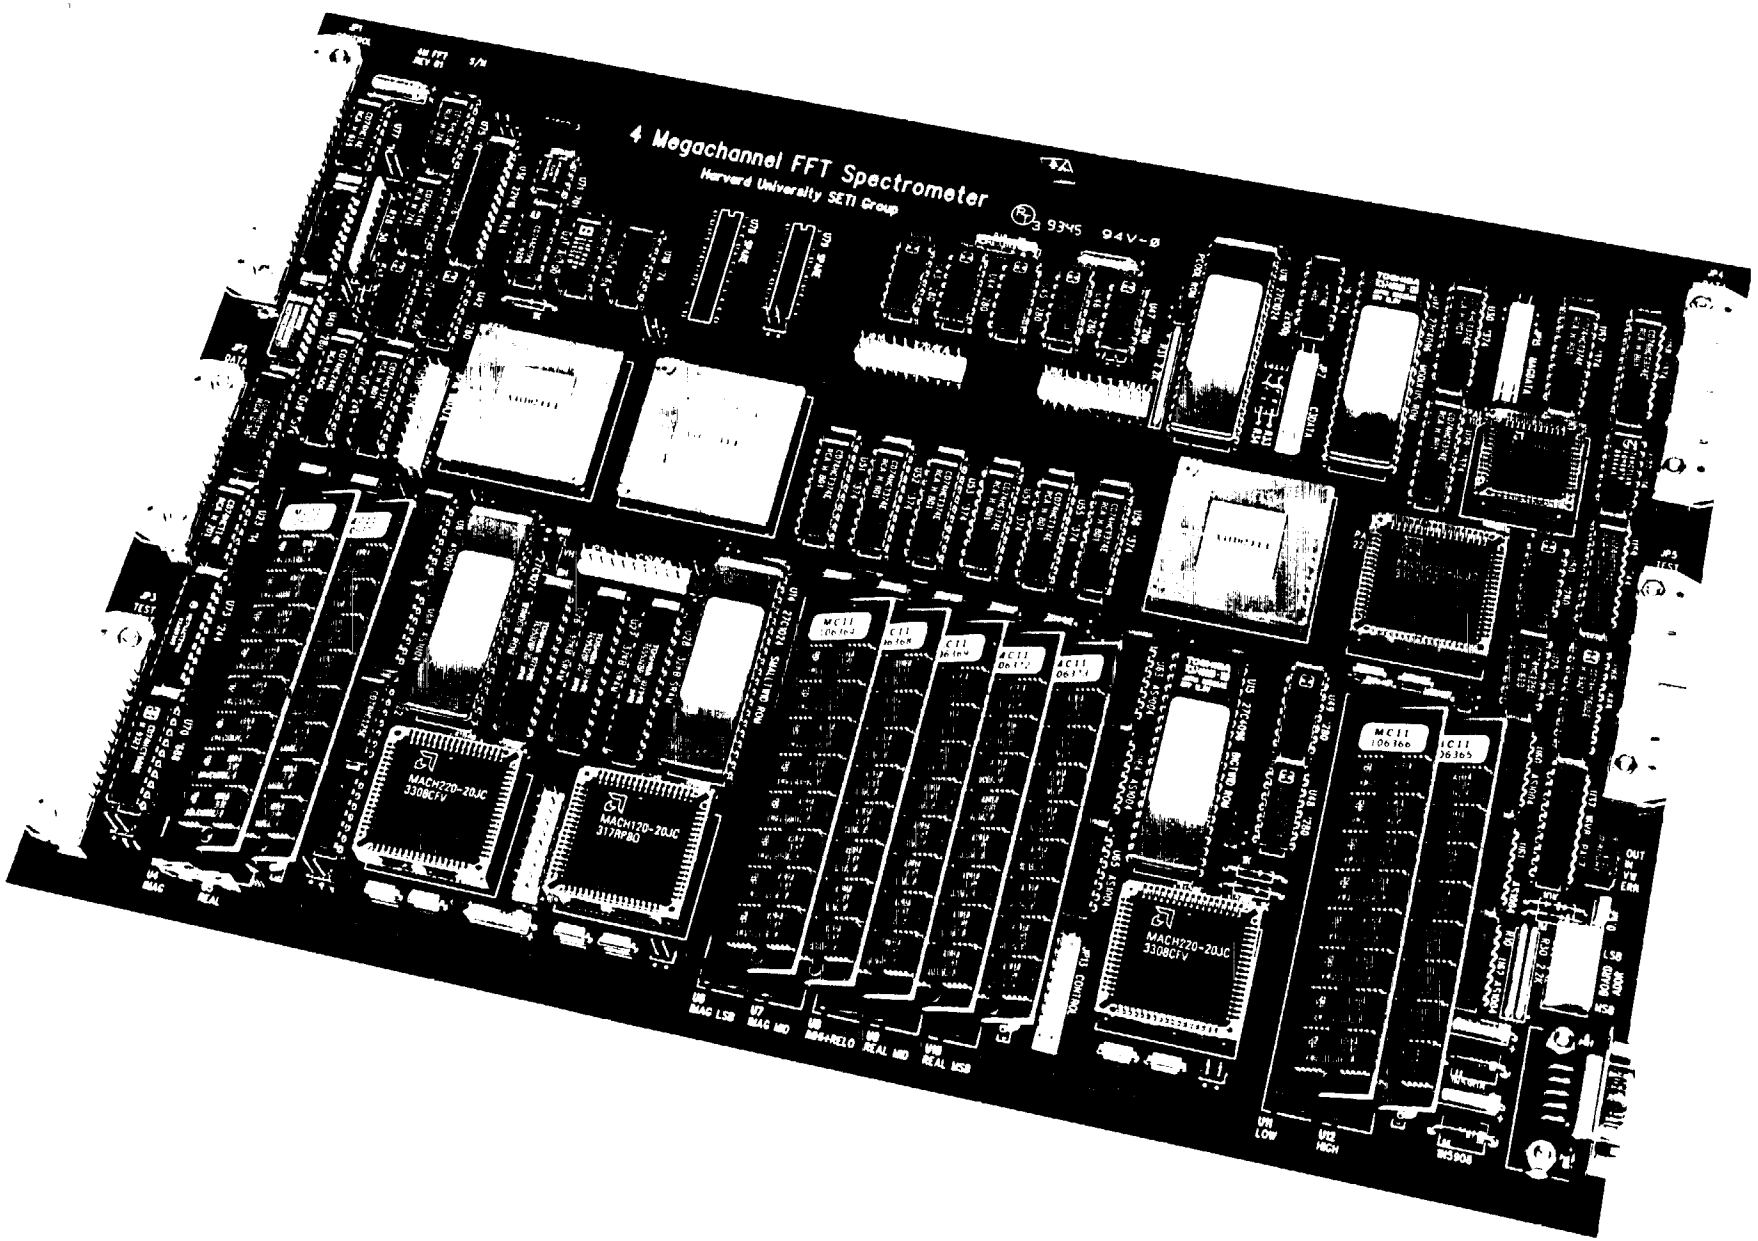

Figure 5. Finished printed circuit board of the 4-megachannel FFT. It measures 9.2 by 13 inches, and includes 36 MByte of parity DRAM, three FFT ASICs, five complex PLDs, and assorted "glue" logic. The complete BETA system requires 63 boards for its 250 million channels.

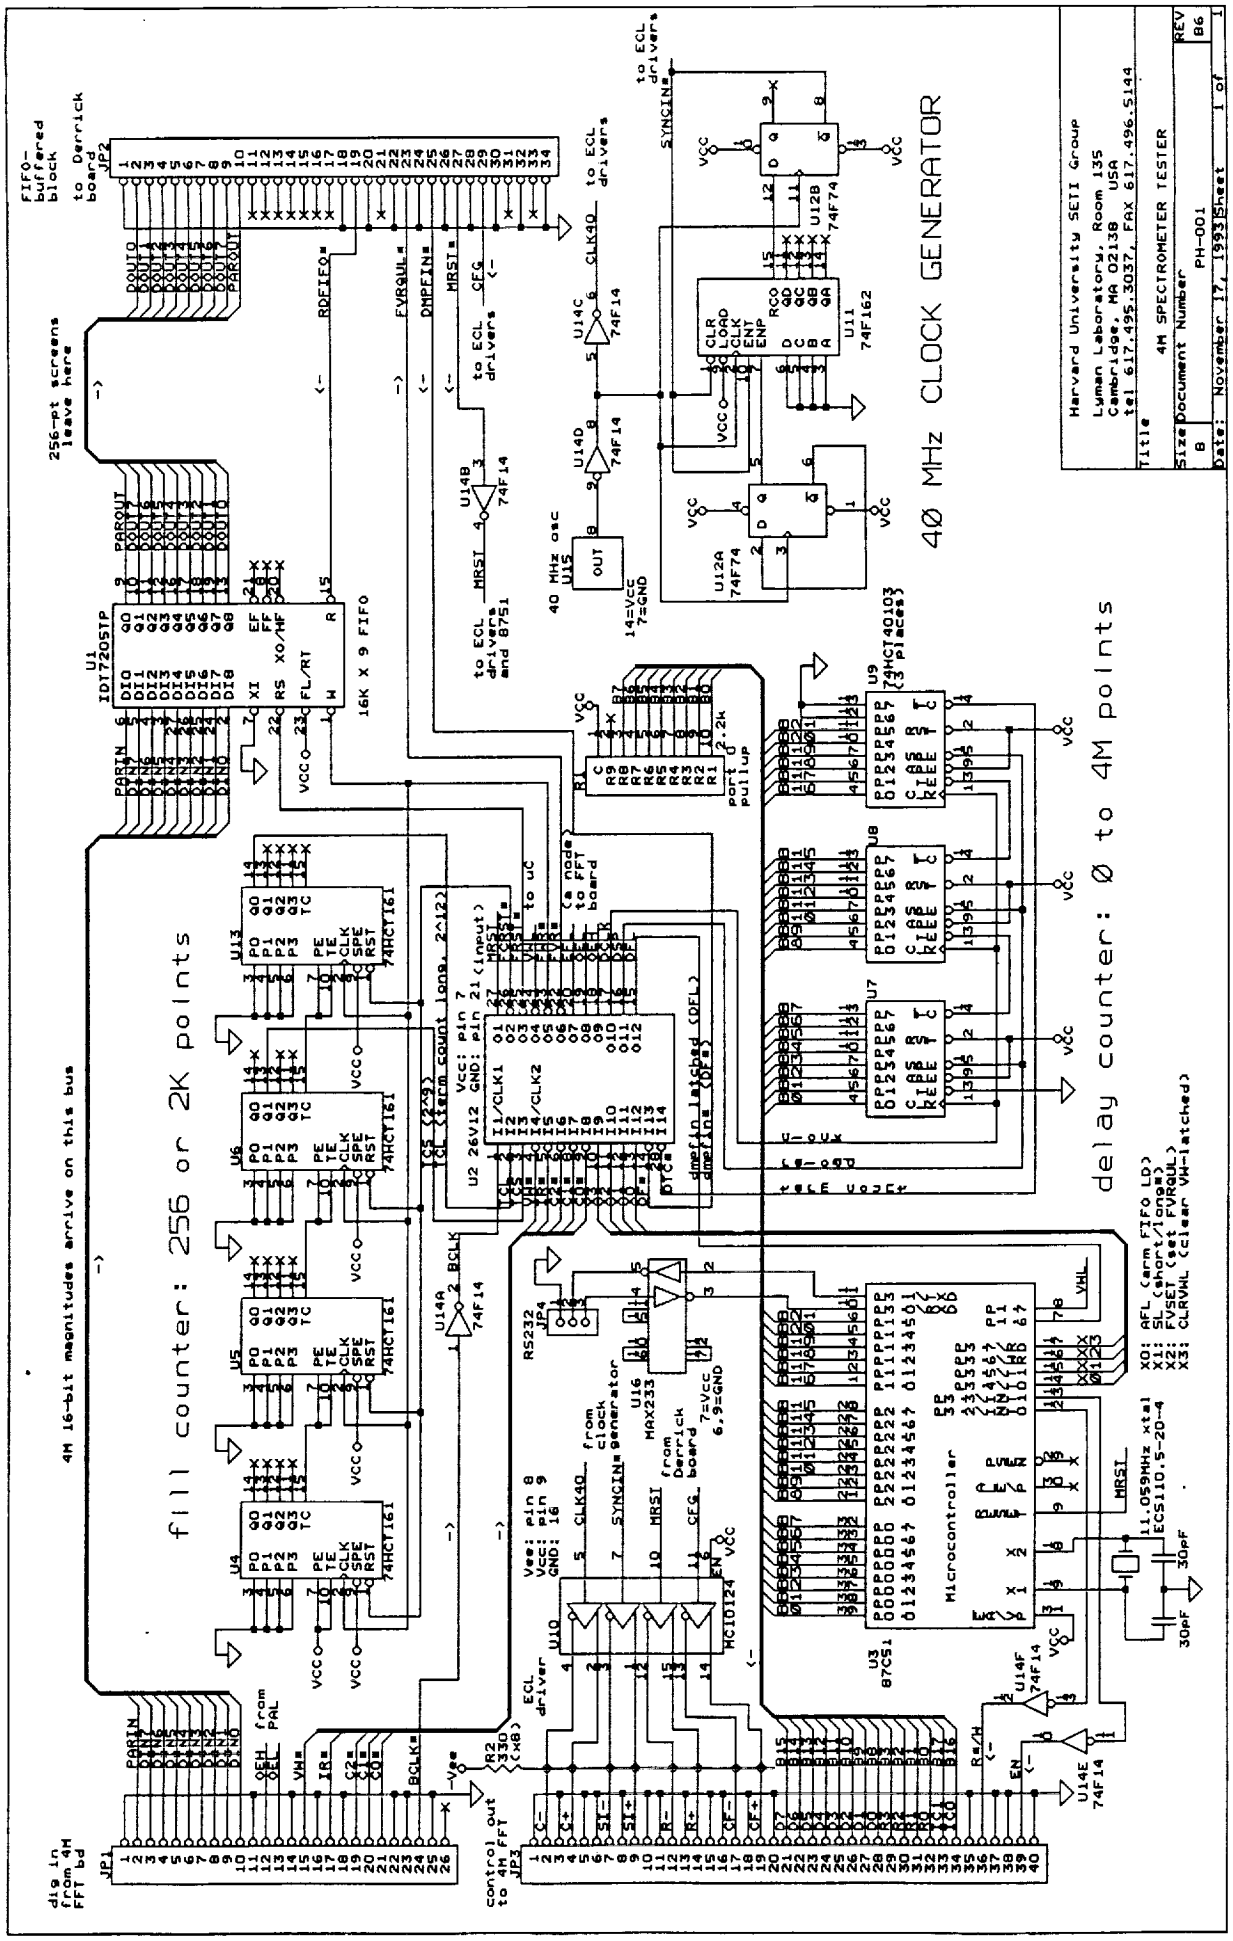

Figure 6. Tester board, to control the FFT board and grab 256-point spectral "snapshots." The latter are either logged to a PC, or plotted to a vector display by the circuit in Figure 7.

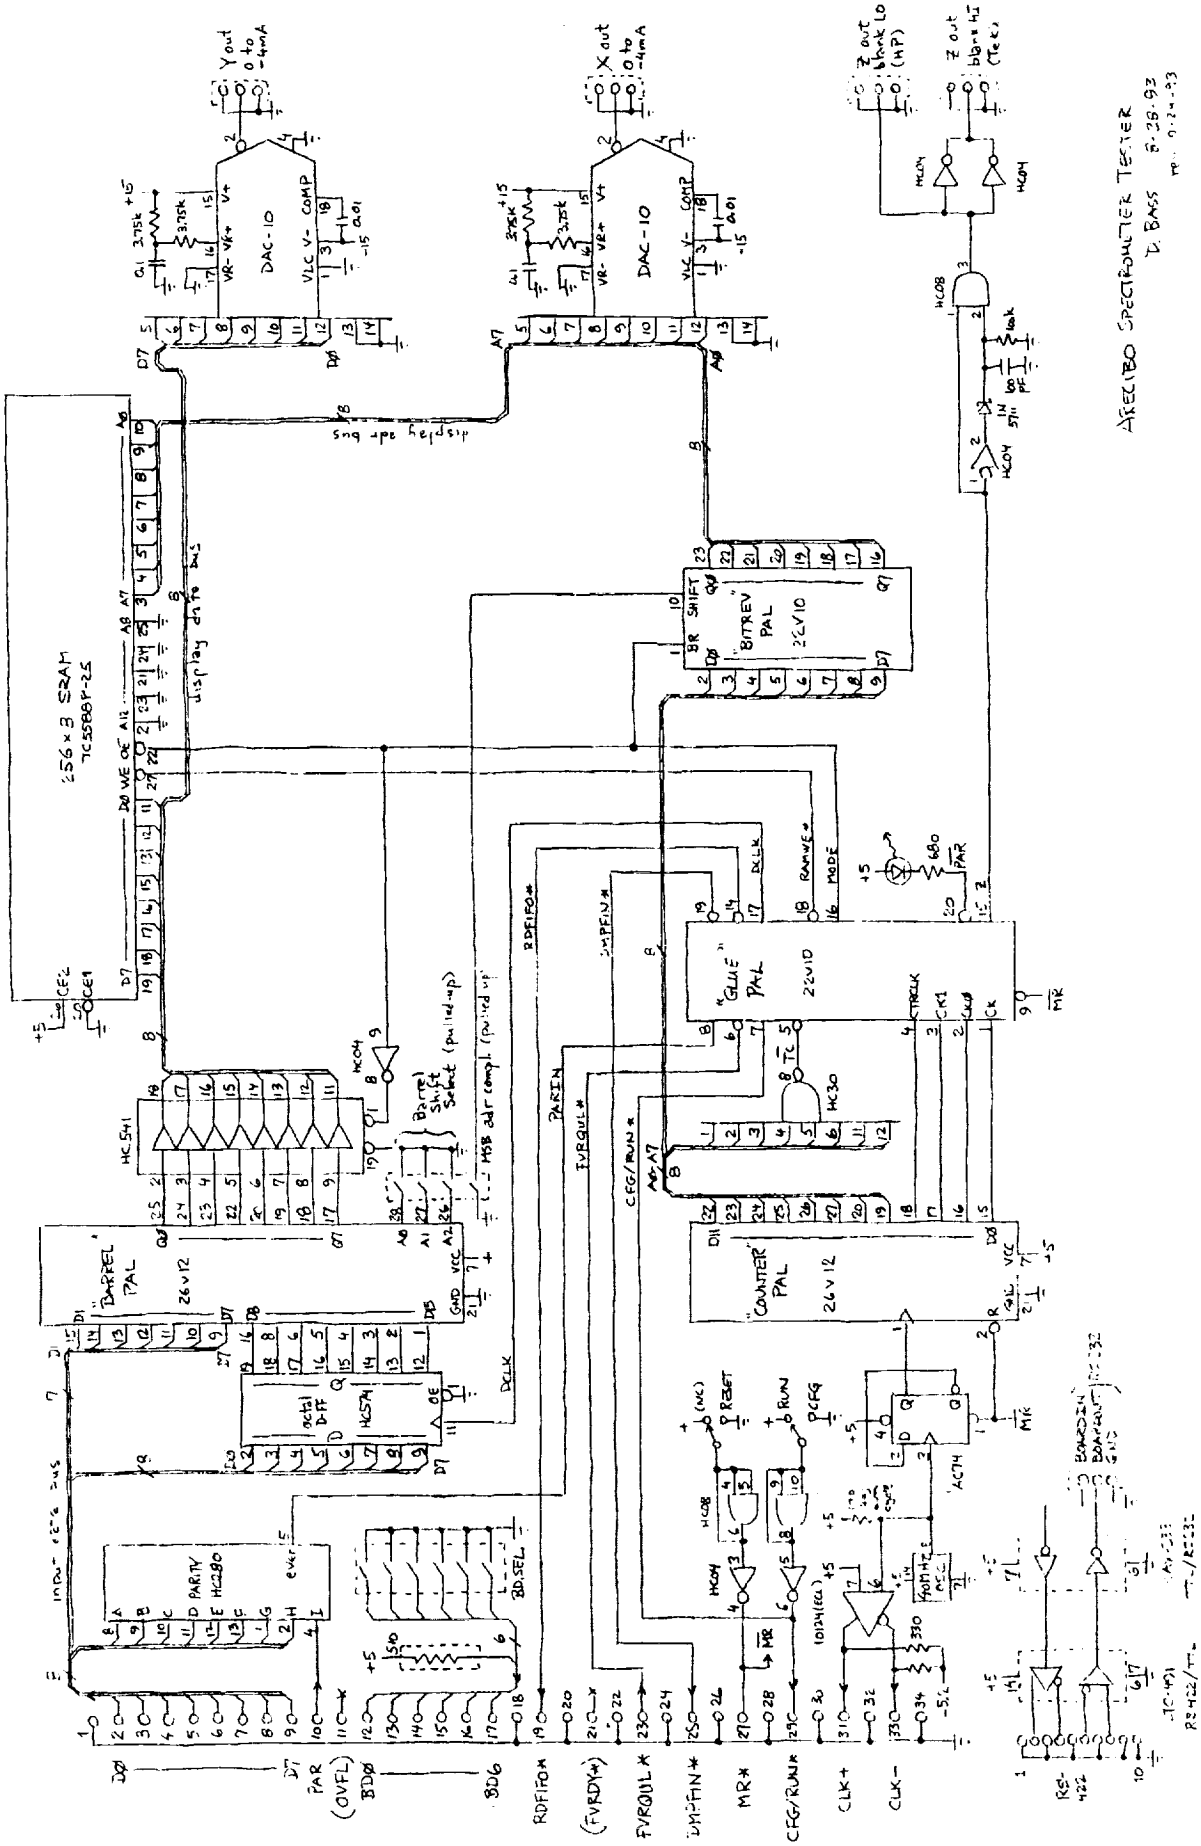

Figure 7. SRAM-based video buffer, for vector display of 256-point spectral snapshots.

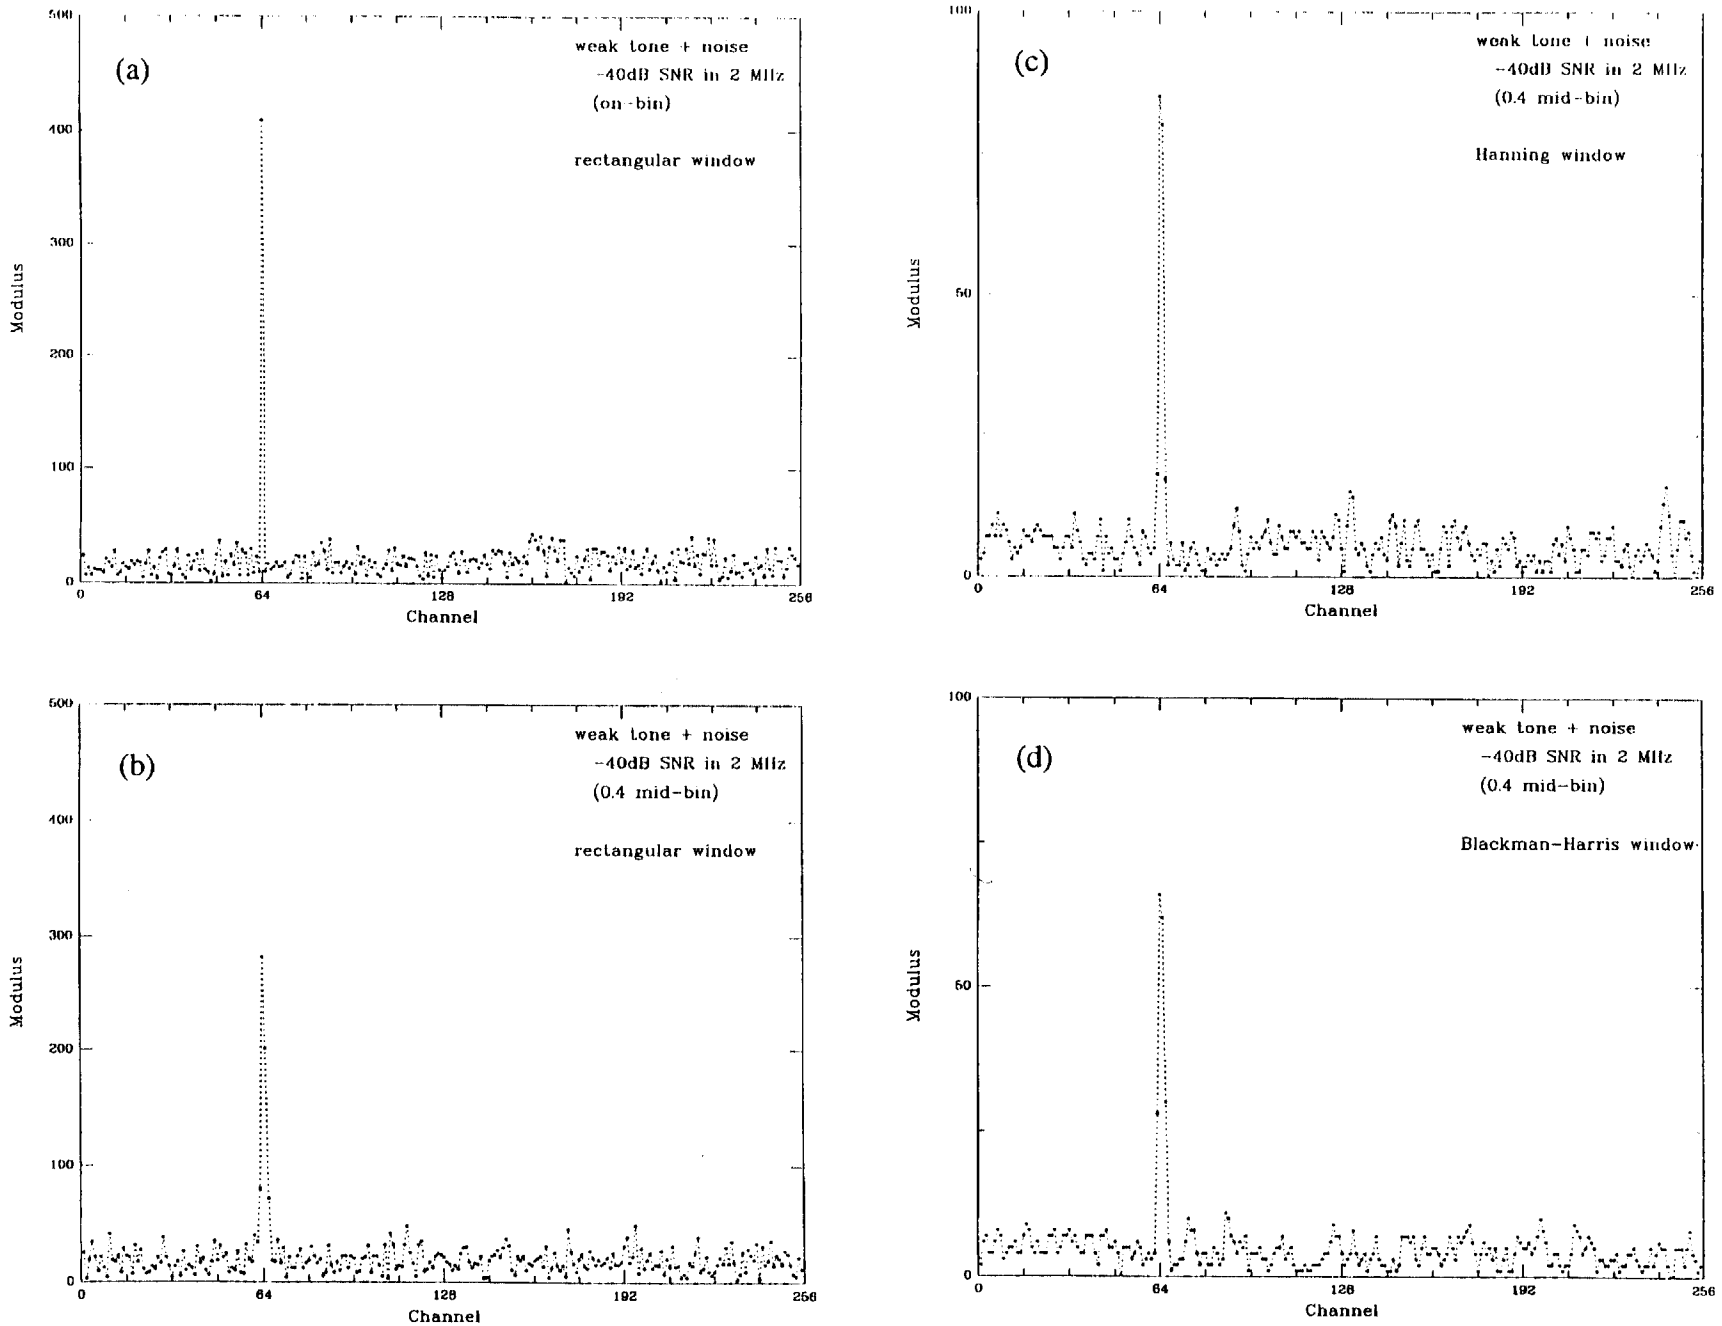

Figure 8. Detection of a weak tone in broadband noise; the appropriate 256-point segment from the full 4M points is displayed. The vertical axis is linear in modulus. (a) -40dBc on-bin, no window -- maximum S/N ratio, but unrealistic; (b) -40dBc between bins, no window; (c) same, Hanning window; (d) same, Blackman-Harris window. The effect of windows is minor in this case of modest single-tone SNR; see Figure 9 for the real significance of windows.

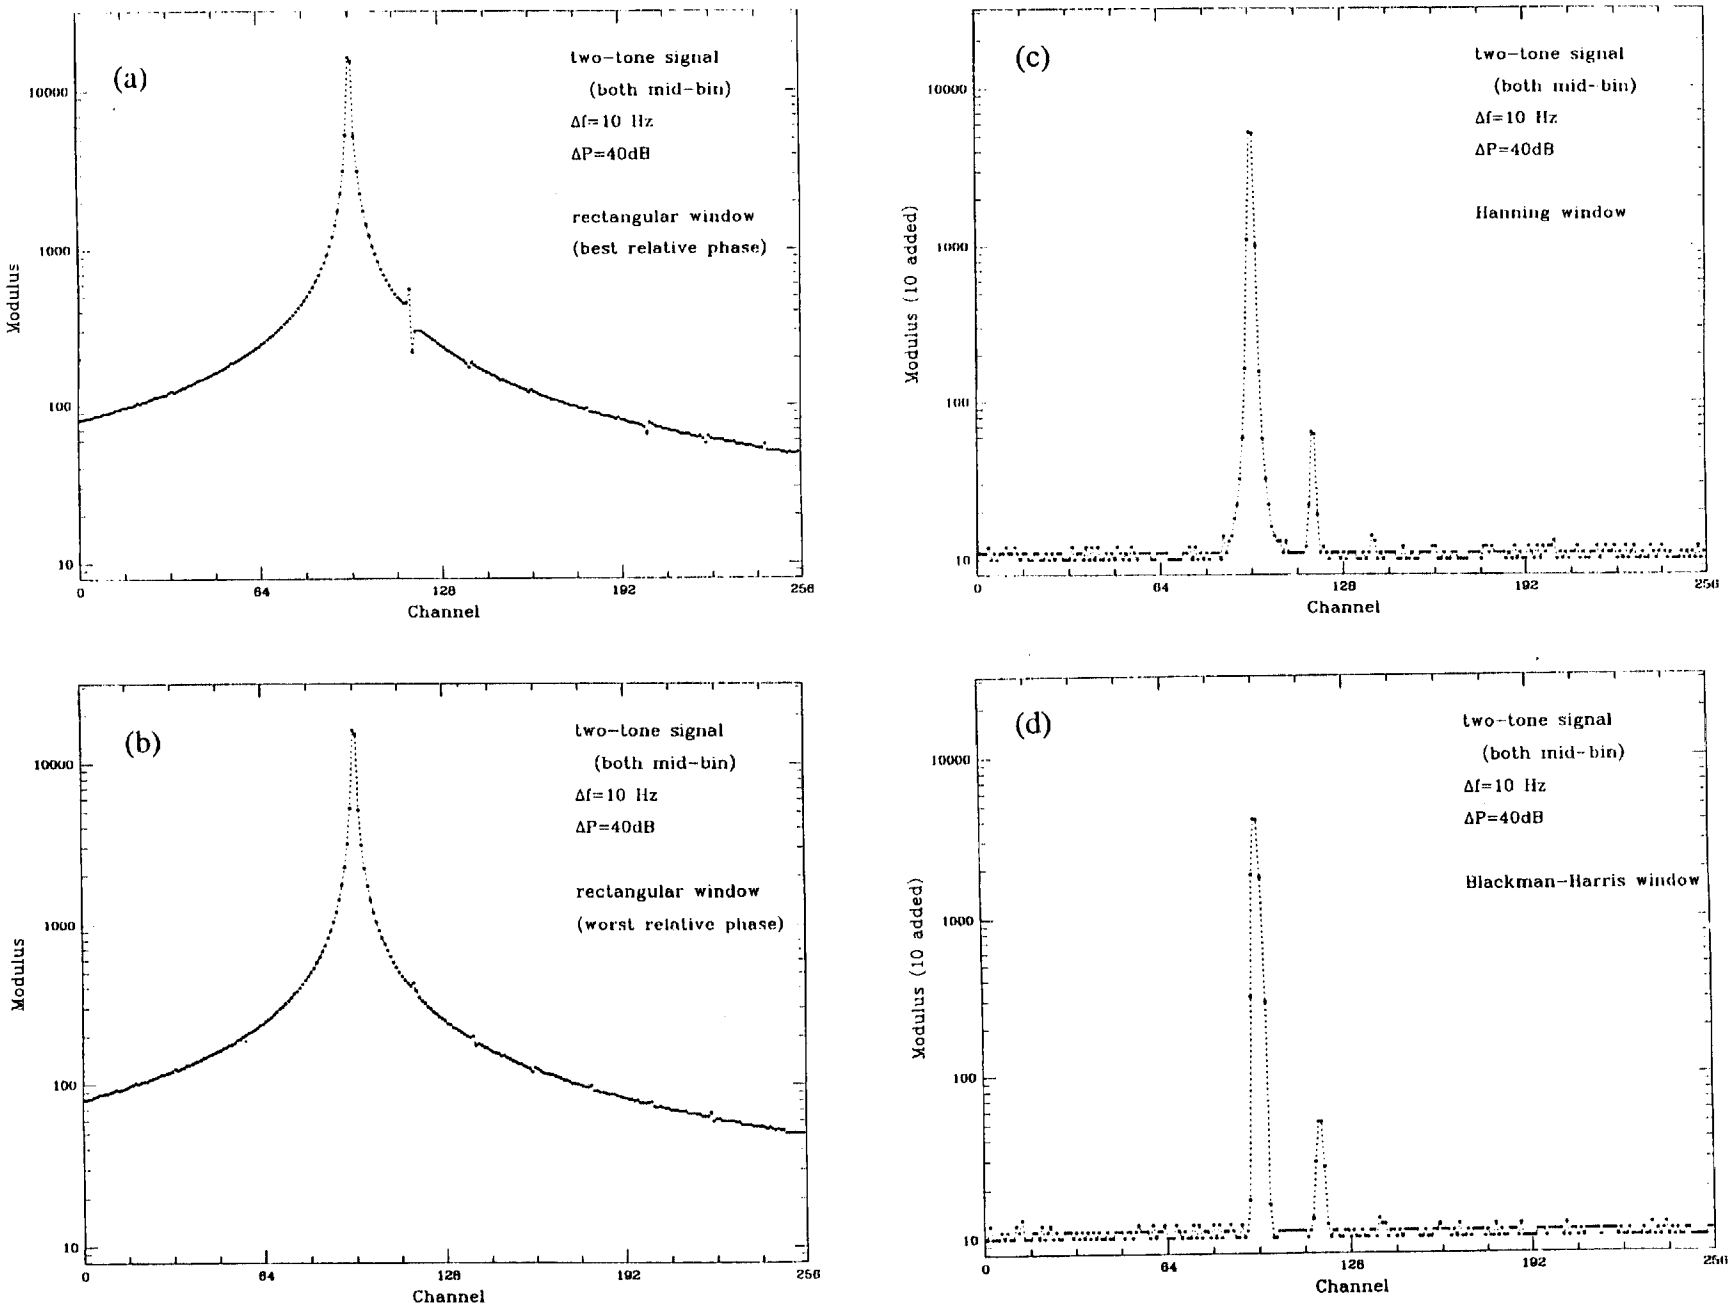

Figure 9. Two-tone signal detection, separated 10 Hz in frequency and 40dB in power; both tones are placed between bins, the most difficult case. The vertical axis is logarithmic in modulus; the spectra in (c) and (d) have 10 added to the output, to suppress the visual raggedness that otherwise results from logarithmic plots of small numbers. (a) and (b) no window, examples of best and worst phasing of the weaker signal; (c) and (d) same, but windowed by Hanning and Blackman-Harris. The reduction of spectral "leakage" is stunning.

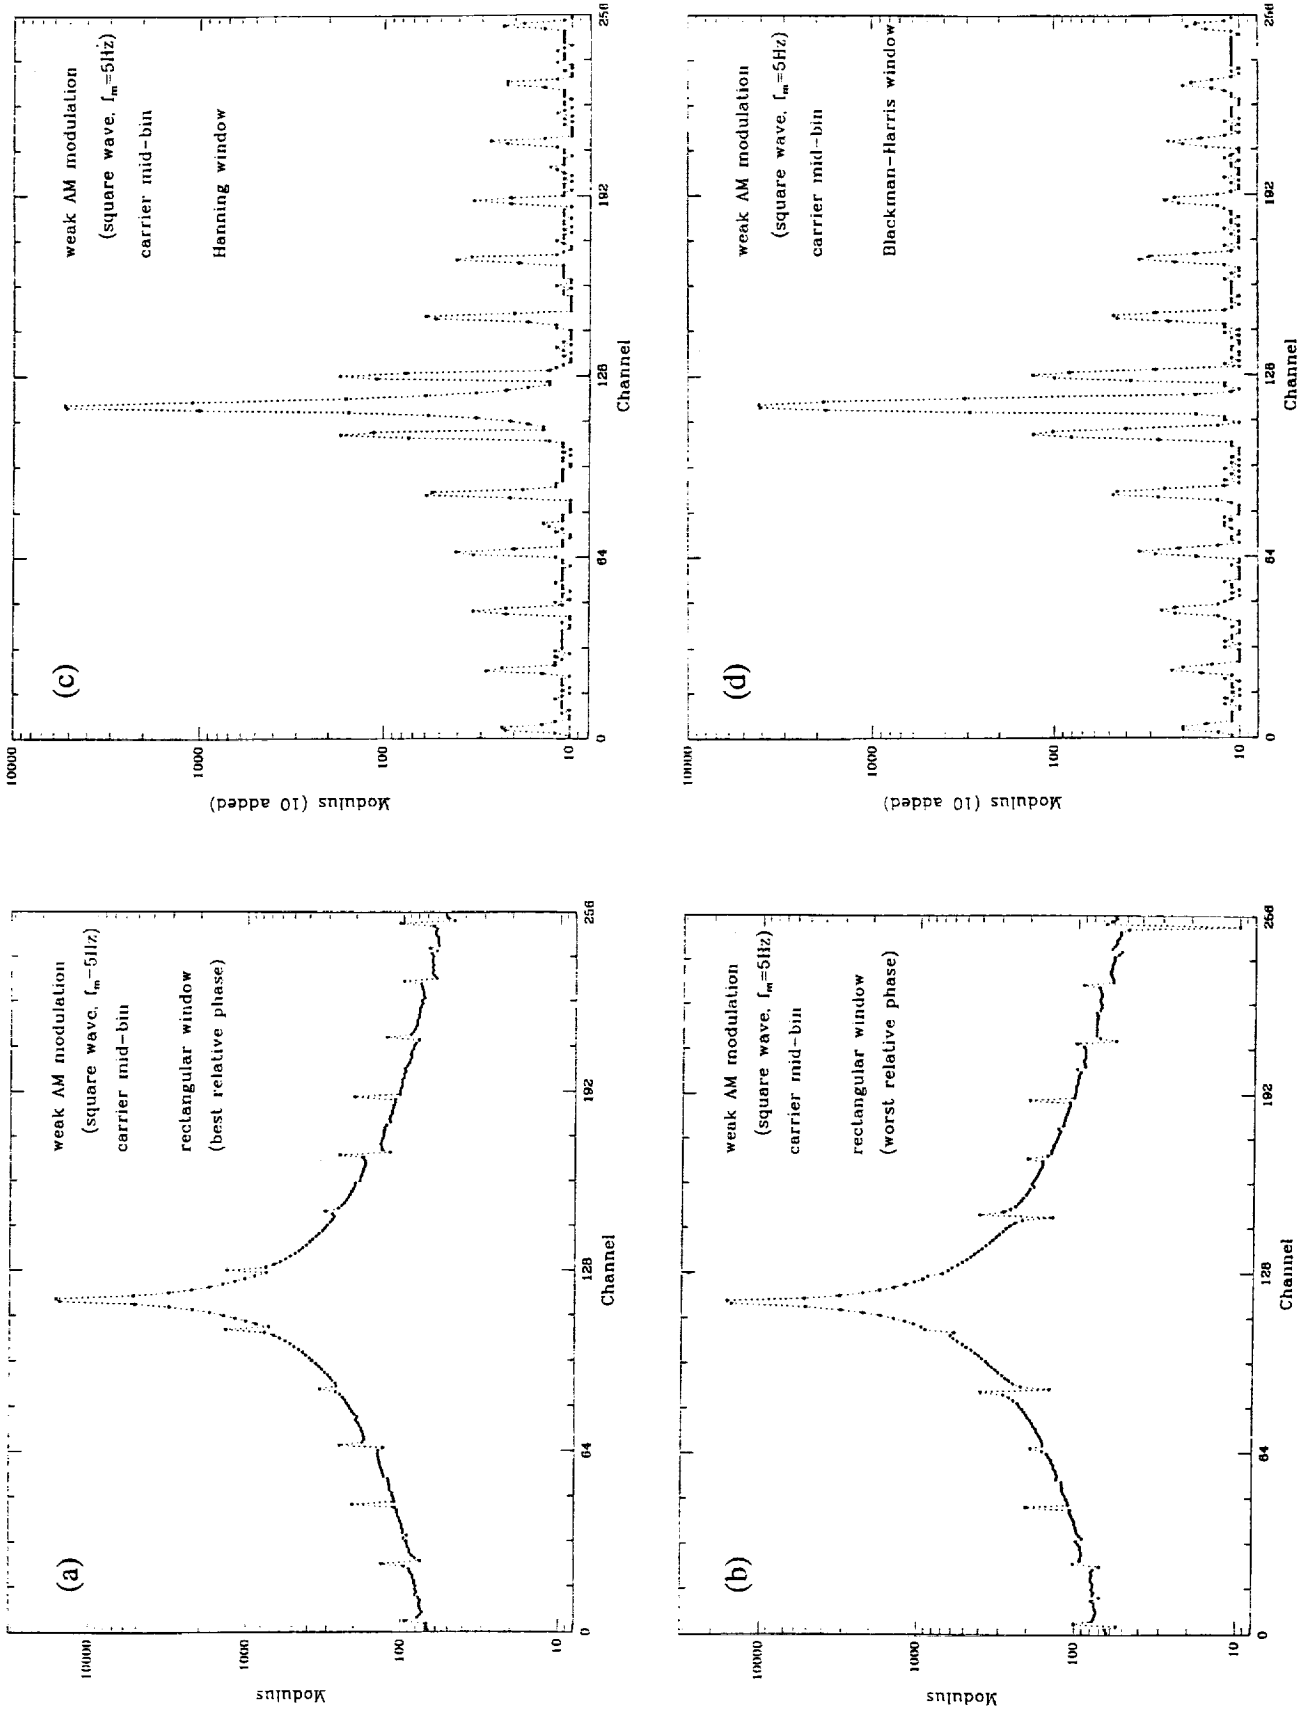

Figure 10. A carrier, placed between bins, weakly AM modulated by a 5 Hz square wave. The four plots are as in Figure 9. Once again, windowing of the time series is essential for serious spectral estimation.

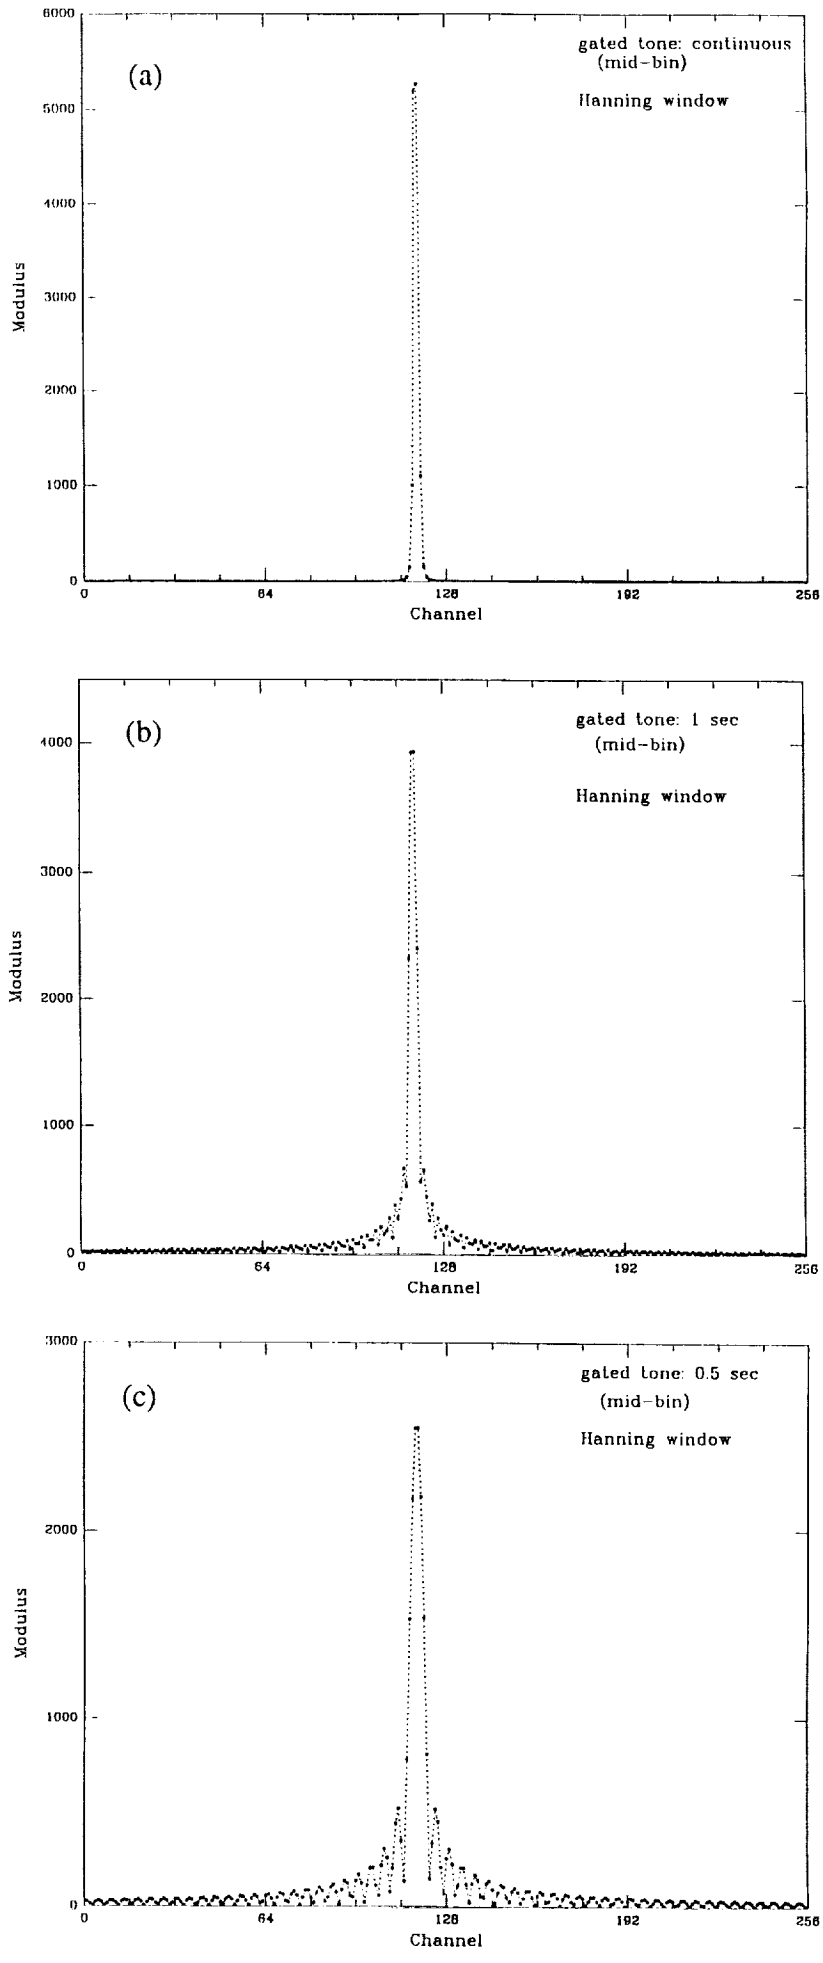

Figure 11. Tone bursts of a carrier placed between bins. (a) continuous signal, 2 seconds in duration; (b) 1-second burst; (c) 0.5-second burst.

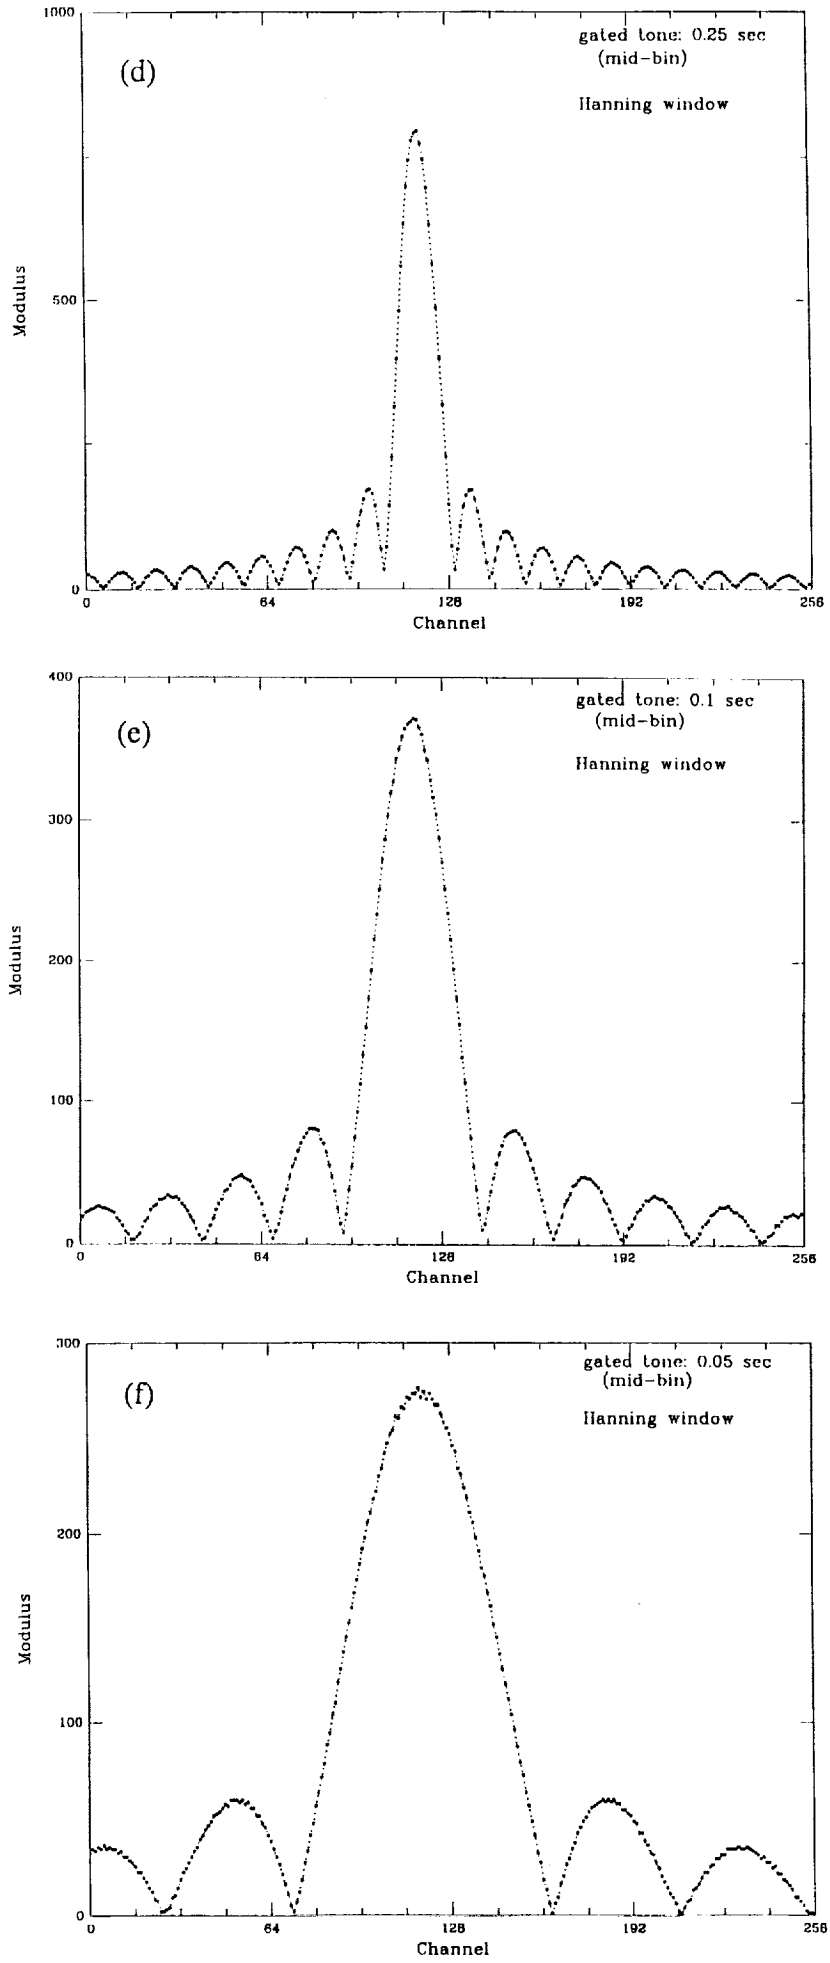

Figure 11 (continued). (d) 0.25-second burst; (e) 0.1-second burst; (f) 0.05-second burst. Note the successive broadening by a *sinc* envelope, and corresponding reduction of peak spectral energy. A short burst of signal is hard to confuse with a continuous carrier.

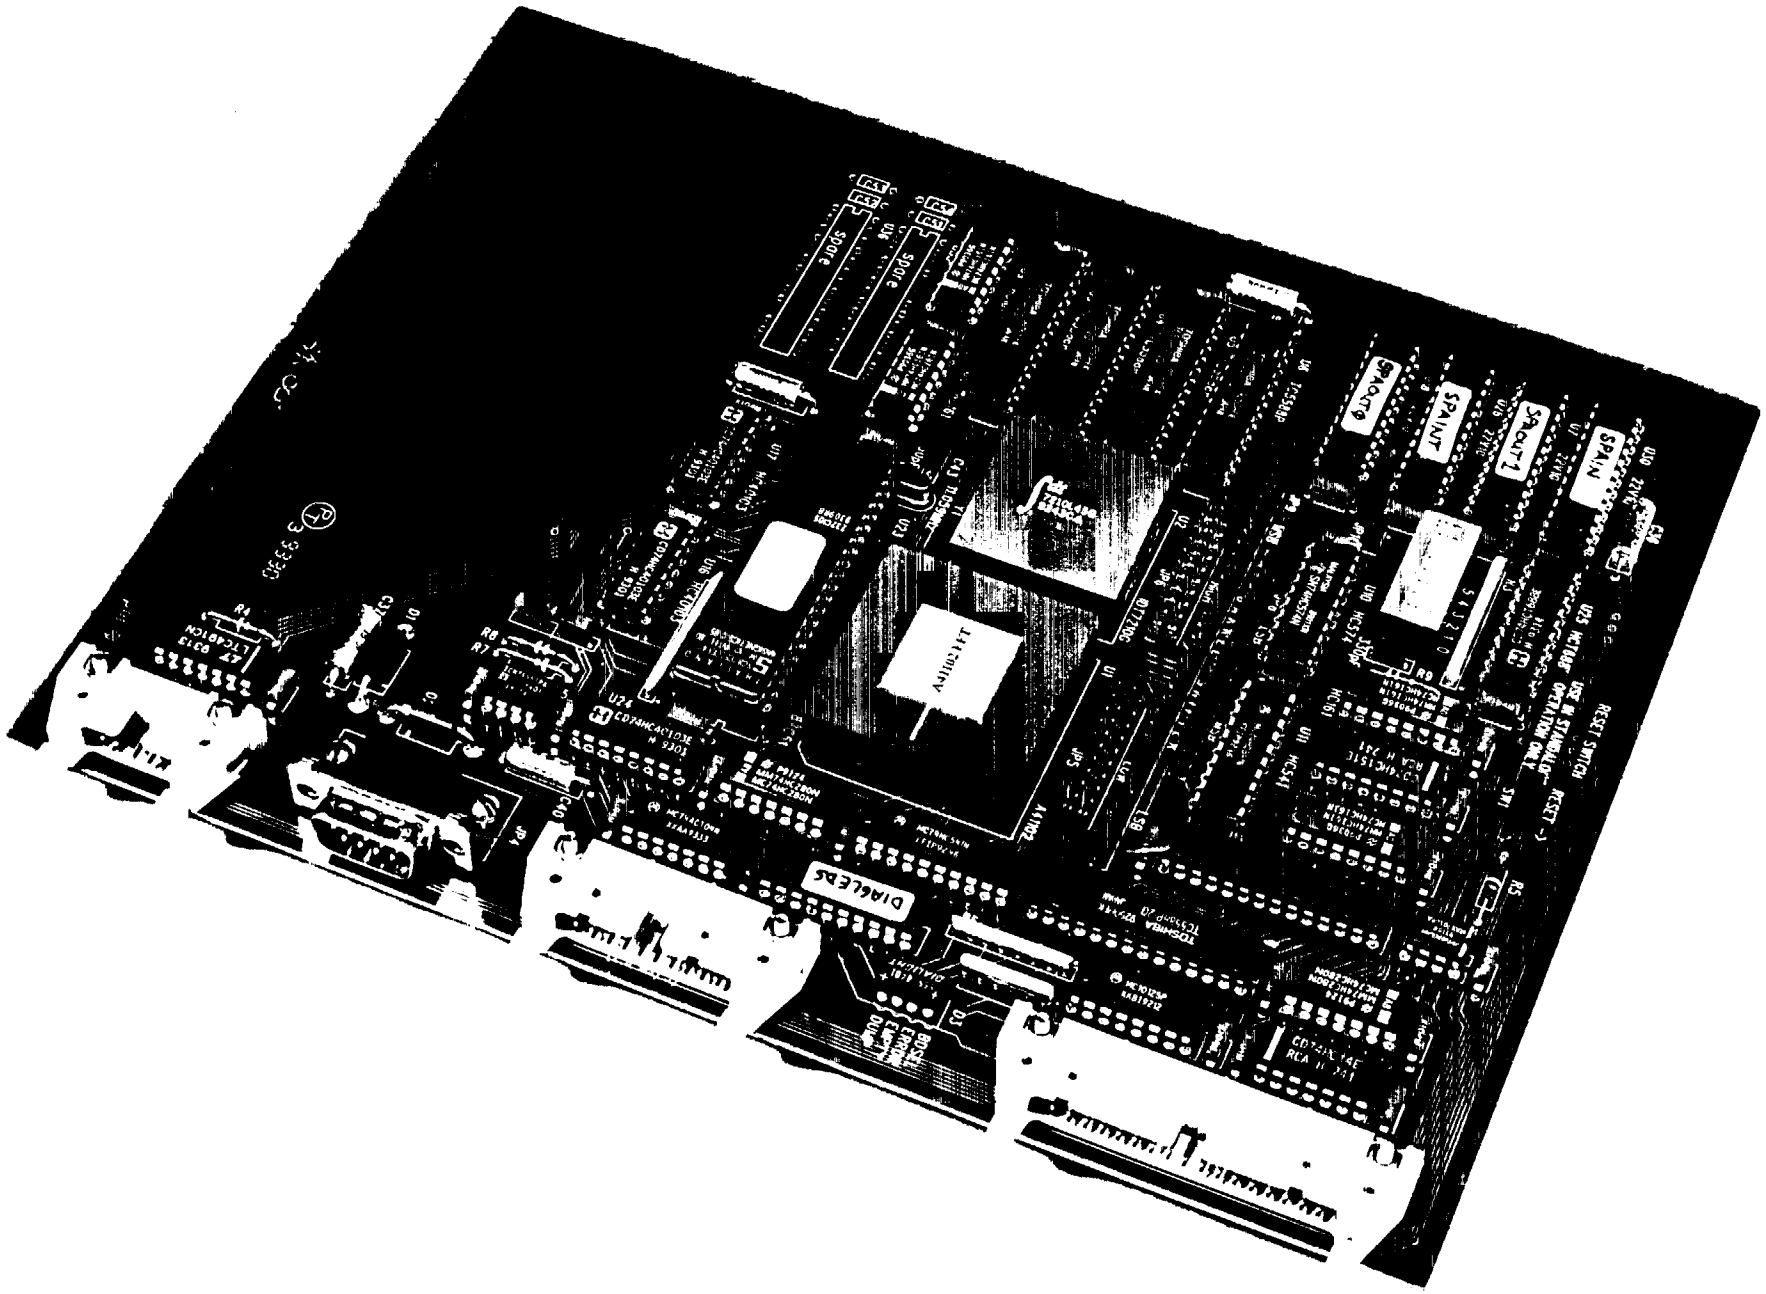

Figure 12. Radioastronomy spectrometer with power accumulator (SPA). A single FFT chip computes 256-point spectra with 2.5 MHz instantaneous bandwidth, at the rate of 10,000 spectra per second. The other prominent chip is the multiplier-accumulator, which computes and accumulates square modulus. The board supports full-speed spectral dumps, or accumulations of any number of successive spectra up to 30 seconds. The 4-layer board uses 12/12 single-track design rules.

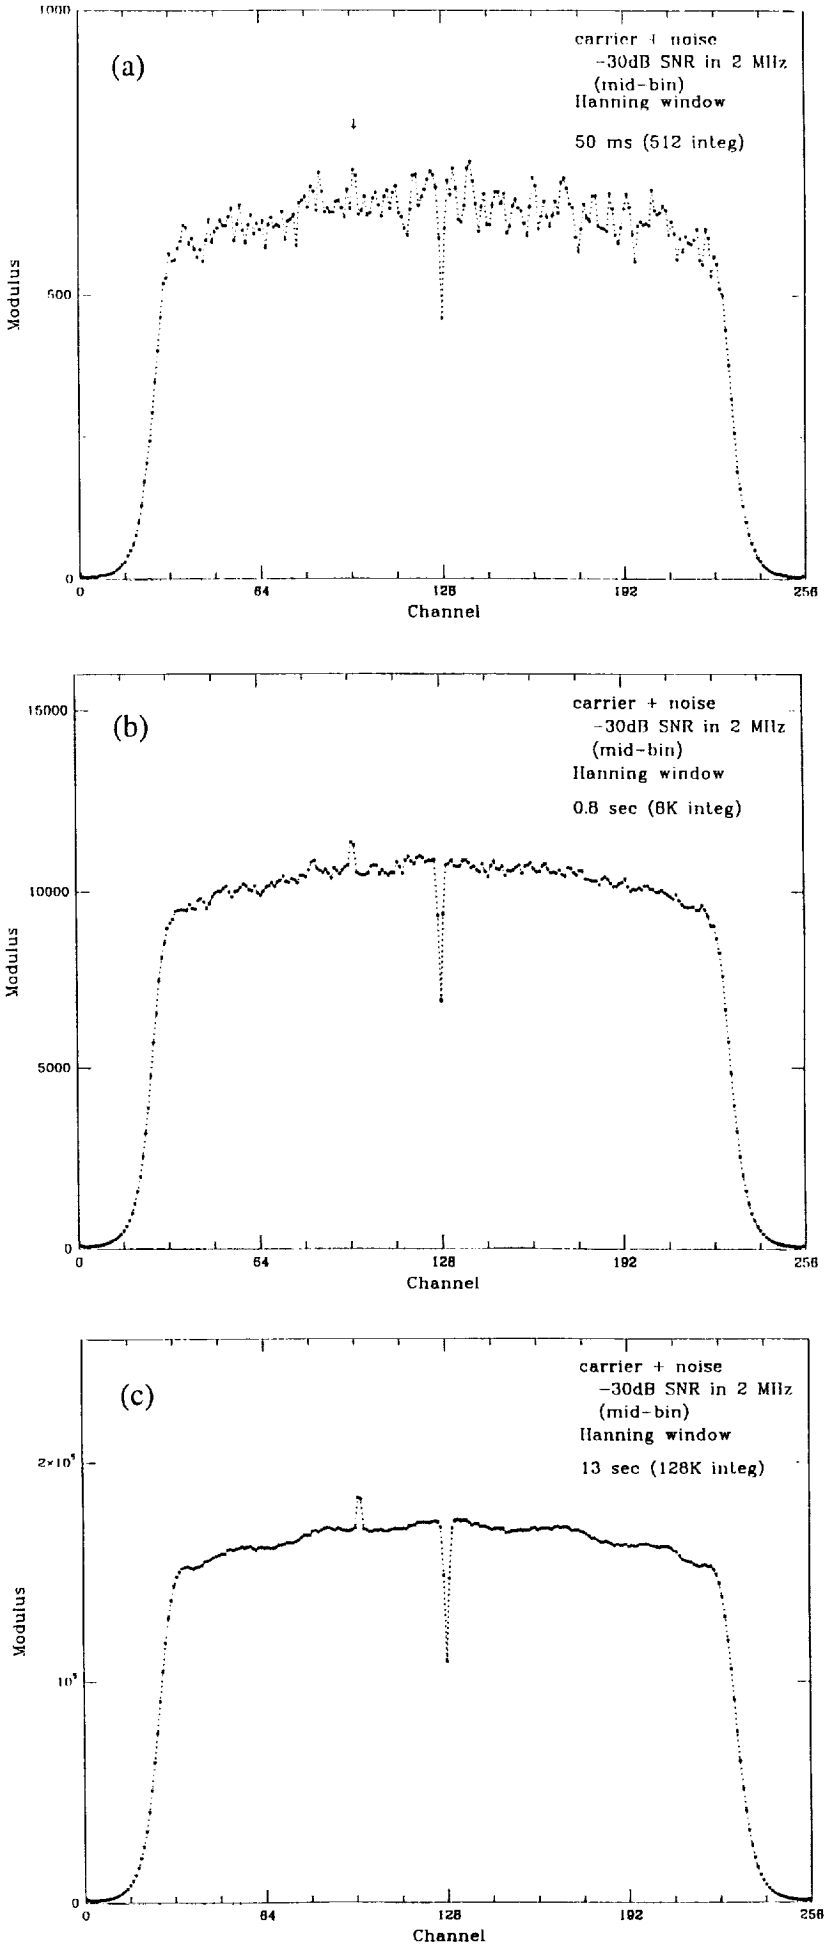

Figure 13. Accumulated 256-point spectra of a weak carrier in broadband noise. (a) 512 spectra (50 ms integration); (b) 8K spectra (0.8 second integration); (c) 128K spectra (13 second integration). The FFT oversamples the 2 MHz filtered baseband, evident as band-edge rolloff; the dip at DC is due to AC coupling at the digitizer input. Ripple in the anti-alias LPF response can be seen in the well-averaged spectrum in (c).

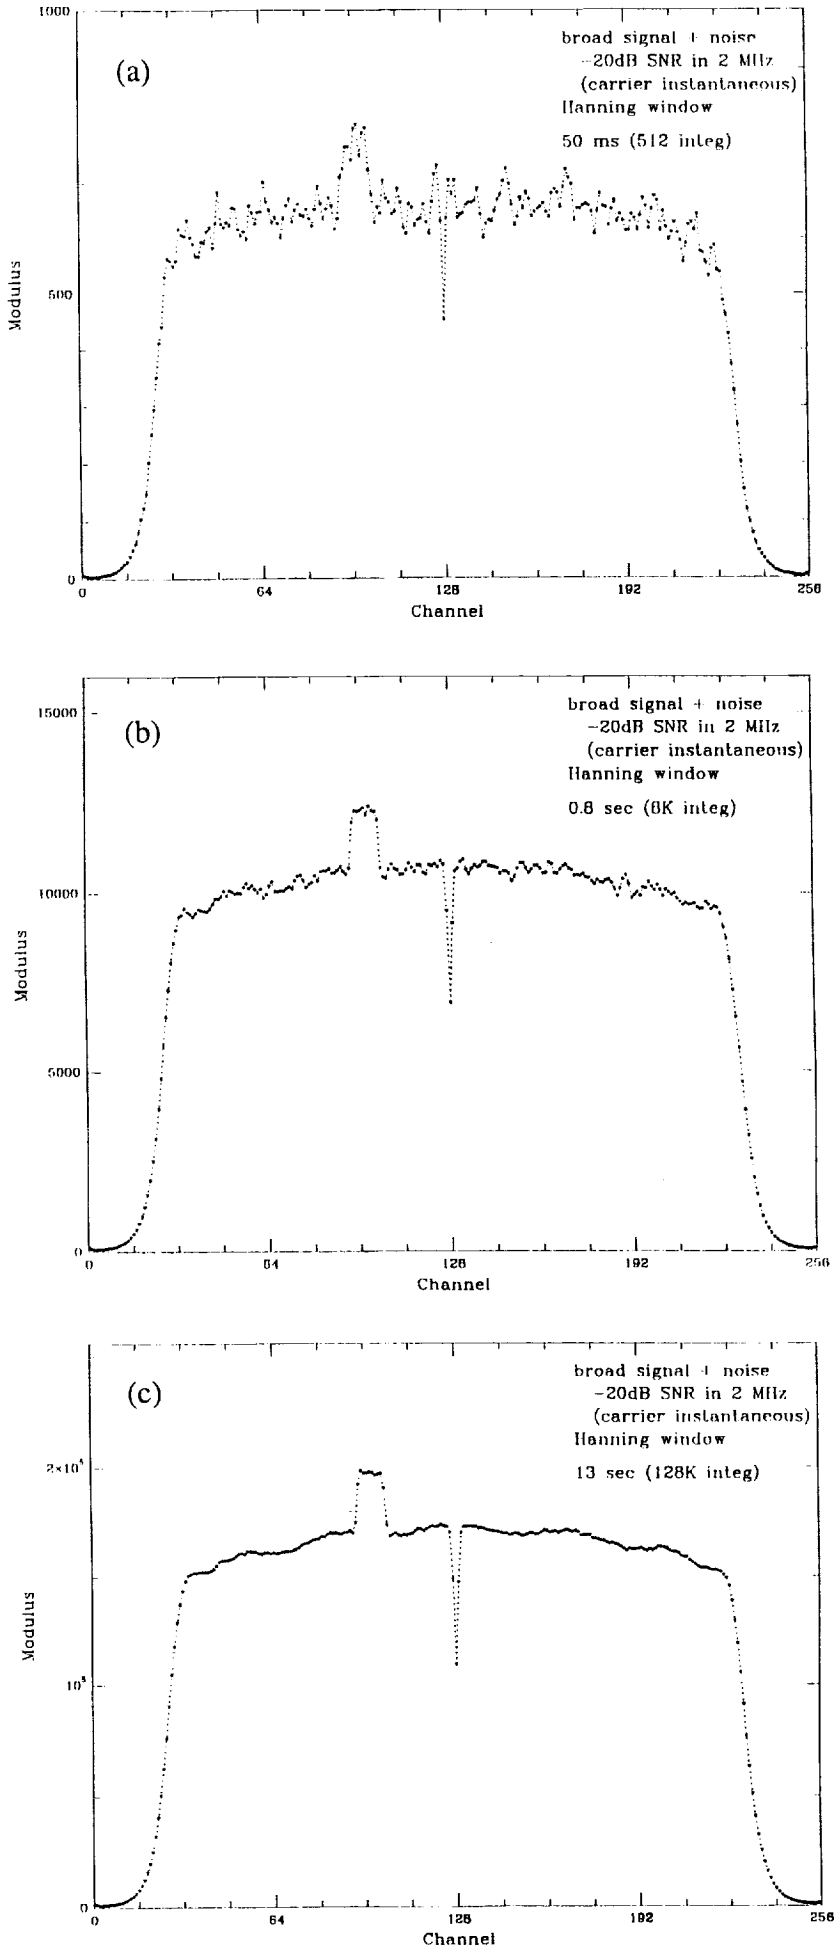

Figure 14. Accumulated 256-point spectra of a flat-topped spectral feature. The 100kHz wide feature was produced by triangle-wave FM. Parameters same as in Figure 13.